Sun™ LSI 106x RAID User's Guide

Sun Microsystems, Inc.
www.sun.com

Part No. 820-4933-12
August 2008

Submit comments about this document at: `http://www.sun.com/hwdocs/feedback`

Copyright © 2008 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Unpublished - rights reserved under the Copyright Laws of the United States.

THIS PRODUCT CONTAINS CONFIDENTIAL INFORMATION AND TRADE SECRETS OF SUN MICROSYSTEMS, INC. USE, DISCLOSURE OR REPRODUCTION IS PROHIBITED WITHOUT THE PRIOR EXPRESS WRITTEN PERMISSION OF SUN MICROSYSTEMS, INC.

This distribution may include materials developed by third parties.

Sun, Sun Microsystems, the Sun logo, Java, Solaris, Sun Fire 4140, Sun Fire 4240, and Sun Fire 4440 are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

AMD Opteron and Opteron are trademarks of Advanced Micro Devices, Inc. Intel is a registered trademark of Intel Corporation.

This product is covered and controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

Use of any spare or replacement CPUs is limited to repair or one-for-one replacement of CPUs in products exported in compliance with U.S. export laws. Use of CPUs as product upgrades unless authorized by the U.S. Government is strictly prohibited.

Copyright © 2008 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

Non publie - droits réservés selon la législation des Etats-Unis sur le droit d'auteur.

CE PRODUIT CONTIENT DES INFORMATIONS CONFIDENTIELLES ET DES SECRETS COMMERCIAUX DE SUN MICROSYSTEMS, INC. SON UTILISATION, SA DIVULGATION ET SA REPRODUCTION SONT INTERDITES SANS L AUTORISATION EXPRESSE, ECRITE ET PREALABLE DE SUN MICROSYSTEMS, INC.

Cette distribution peut inclure des éléments développés par des tiers.

Sun, Sun Microsystems, le logo Sun, Java, Solaris et Sun Fire X4140, Sun Fire X4240 et Sun Fire X4440 sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

AMD Opteron et Opteron sont marques déposées de Advanced Micro Devices, Inc. Intel est une marque déposée de Intel Corporation

Ce produit est soumis à la législation américaine sur le contrôle des exportations et peut être soumis à la règlementation en vigueur dans d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes biologiques et chimiques ou du nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou reexportations vers les pays sous embargo américain, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de maniere non exhaustive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine sur le contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

L'utilisation de pièces détachées ou d'unités centrales de remplacement est limitée aux réparations ou à l'échange standard d'unités centrales pour les produits exportés, conformément à la législation américaine en matière d'exportation. Sauf autorisation par les autorités des Etats-Unis, l'utilisation d'unités centrales pour procéder à des mises à jour de produits est rigoureusement interdite.

Please
Recycle

Contents

Preface xiii

Part I BIOS RAID Configuration Utility

1. Introduction to Integrated RAID 1

Introduction 1

Integrated RAID Features 2

Using this Manual 2

2. Overview of Integrated Mirroring and Integrated Mirroring Enhanced 3

Introduction 3

IM and IME Features 4

IM/IME Description 5

Integrated RAID Firmware 7

Re-synchronization with Concurrent Host I/O Operation 7

Metadata Support 7

Hot Swapping 7

SMART Support 8

Hot Spare Disk 8

Media Verification 8

Disk Write Caching 8

Write Journaling 8

Fusion-MPT Support 9

3. Creating IM and IME Volumes 11

IM/IME Configuration Overview 11

Creating IM and IME Volumes 12

- To Create an IM Volume 12
- To Create an IME Volume 15

Creating a Second IM or IME Volume 15

- To Create a Second IM or IME Volume Method I 15
- To Create a Second IM or IME Volume Method II 16

Managing Hot Spares 16

- To Create Global Hot-Spare Disks 16

Other Configuration Tasks 18

- To View Volume Properties 18

Synchronizing an Array 19

- To Synchronize an Array 19

Activating an Array 19

- To Activate an Array 19

Deleting an Array 20

- To Delete an Array 20

Locating a Disk Drive or Multiple Disk Drives in a Volume 20

Selecting a Boot Disk 21

- To Select a Boot Disk 21

4. Overview of Integrated Striping 23

Introduction 23

IS Features 24

IS Description 24

Integrated Striping Firmware 26

Metadata Support 26

SMART Support 26

Disk Write Caching 26

Fusion-MPT Support 26

5. Creating Integrated Striping Volumes 27

IS Configuration Overview 27

Creating IS Volumes 28

▼ To Create IS Volumes 28

Creating a Second IS Volume 30

▼ To Create a Second IS Volume, Method I 31

▼ To Create a Second IS Volume, Method II 31

Other Configuration Tasks 31

▼ To View IS Volume Properties 31

Activating an Array 32

▼ To Activate a Selected Array 32

▼ To Delete an Array 32

Locating a Disk Drive or Multiple Disk Drives in a Volume 33

Selecting a Boot Disk 33

▼ To Select a Boot Disk 33

Part II LSI `cfggen` Utility

6. The LSI `cfggen` Utility 37

Installing the `cfggen` Utility 37

Overview of `cfggen` 38

`cfggen` Syntax 38

`cfggen` Command Conventions 39

Common Command-Line Parameters 39

Supported Commands 40

`auto` Command 40

`create` Command 42

`delete` Command 44

`display` Command 44

`hotspare` command 48

`list` command 49

`rebuild` command 50

`status` Command 51

Monitoring and Managing RAID Arrays 53

- To Create a RAID 0 Array 53
- To Fail RAID 0 55
- To Create a RAID 1 Array 56
- To Rebuild a RAID 1 Array 56
- To Delete a RAID Array 57

Part III MegaRAID Storage Manager

7. MegaRAID Storage Manager (MSM) Installation 61

Overview 61

Installing the Program 62

Installing MSM on the Windows OS 62

- To Install MSM on the Windows OS 62

Windows Installation Error Messages 63

Installing Windows Drivers 63

- To Install MSM on the Linux OS 63

- To Create a Partition and Filesystem on Linux OS 65

8. Using MegaRAID Storage Manager 67

Starting the MSM Program 67

▼ To Start MSM on the Windows 2003 Server 68
▼ To Start MSM on a RHEL 4 Server 68
▼ To Start MSM on a SLES 10 Server 68
Running MSM 68
▼ To Access Servers on Alternate Subnets 70
▼ To Log in to MSM 70
MSM Windows 74
Physical/Logical View Panel 74
Event Log Panel 75
Properties/Operations/Graphical View Panel 76
Dual-Path SAS Disk Drives 77
Menu Bar 77
Using the MSM RAID Configuration Wizard 78
▼ To Start the MSM RAID Configuration Wizard 78
▼ To Create a RAID Partition and File System (Windows 2003) 83
Changing Virtual Disk Properties 87
▼ To Change a Virtual Disk Property 88
Deleting a Virtual Disk 88
▼ To Delete A Virtual Disk 89
Saving a Storage Configuration to Disk 89
Clearing a Storage Configuration From a Controller 90
▼ To Clear a Configuration from a Controller 90
Monitoring System Events and Storage Devices 90
Monitoring System Events 91
Monitoring Controllers 91
▼ To Display All Controller Information 92
Monitoring Disk Drives 92
▼ To Display Complete Disk Drive Information 92

▼ To Display a Graphical View of a Disk Drive 92

Monitoring Virtual Disks 93

▼ To Display a Graphical View of a Virtual Disk 93

Monitoring Rebuilds and Other Processes 94

Maintaining and Managing Storage Configurations 94

▼ To Scan for New Drives 95

Rebuilding a Drive 95

▼ To Rebuild a Drive on a SAS IR System 95

Putting a Drive Offline or Missing 96

▼ To Put a Drive Offline or Missing 96

Known Issues 97

Sun Fire X4100 M2 /X4200 M2 Server Issues 97

(Windows 2003 Server) MSM-IR 1.19 Does Not Reflect Correct Disk Count Information (CR 6514389) 97

Sun Fire X4600 M2 Server Issues 97

MSM-IR 1.19 Does Not Show Disk Removal Status Correctly in a Non-RAID Configuration (CR 6525255) 97

MSM Server and Client Must Be in Same Subnet (6533271) 97

Sun Blade X6240, X6440 Server Issues 98

Locate Virtual Disk Function Does Not Light LEDs on Disks Controlled by Server Blade (CR 6732326) 98

9. LSI SNMP Utility 99

Installing and Configuring SNMP Service 99

▼ To Install the SNMP Service on the Server Side on Windows Server 2003 100

▼ To Configure the SNMP Service on the Server Side on Windows Server 2003 100

Installing and Configuring the SNMP Server on the Server Side on Linux 101

▼ To Install the SNMP Agent 101

Installing LSI SNMP on a Remote Station 101

▼ To Download the LSI SNMP Agent Files 102

▼ To Install SNMP Files on a Remote Station 102

Part IV P A R T IV - Using `raidctl` With Solaris OS

10. Introduction to `raidctl` 107

What is `raidctl`? 107

When to Use `raidctl` 108

Using `raidctl` to Create RAID Volumes 108

Disk Names 108

Obtaining Disk Names in Canonical Format 109

Obtaining Disk Names in `C.ID.L` Format 109

The `raictl -c` Command 110

The `raidctl -C` Command 111

Other Uses For `raidctl` 112

11. The `raidctl` Man Page 113

The Man Page for `raidctl` 113

Part V Glossary and Index

Glossary 127

Index 133

Figures

FIGURE 2-1 Typical Integrated Mirroring Implementation 5

FIGURE 2-2 Integrated Mirroring Volume 6

FIGURE 2-3 Integrated Mirroring Enhanced with Three Disks 6

FIGURE 3-1 Adapter Properties Screen 13

FIGURE 3-2 Create New Array Screen 14

FIGURE 3-3 Manage Array Screen 17

FIGURE 4-1 Integrated Striping Example 25

FIGURE 4-2 Integrated Striping - Logical and Physical Views 25

FIGURE 5-1 Adapter Properties Screen 29

FIGURE 5-2 Create New Array Screen 30

FIGURE 8-1 MSM Select Server Window (Opening Screen) 68

FIGURE 8-2 MSM Select Server Window (Simplified) 70

FIGURE 8-3 MSM User Name & Password Dialog Box 71

FIGURE 8-4 MSM Physical/Logical Window 73

FIGURE 8-5 MSM Window (Events and Icons) 75

FIGURE 8-6 MSM Window (RAID Creation) 78

FIGURE 8-7 MSM Configuration Wizard Dialog Box 79

FIGURE 8-8 MSM Virtual Disk Creation Dialog Box 81

FIGURE 8-9 MSM RAID Finish Screen 82

FIGURE 8-10 MSM Physical Window (New RAID Array) 83

FIGURE 8-11 Windows Disk Manager Initialize Wizard 84
FIGURE 8-12 Windows Disk Manager Initialize Disk 84
FIGURE 8-13 Windows Disk Manager New Partition Wizard 85
FIGURE 8-14 Windows Display of New RAID File System 86

Preface

This *Sun™ LSI 106x RAID User's Guide* contains instructions for creating and maintaining hardware RAID volumes. It applies to all servers (including blades) that have on-board LSI 106x chips or that use PCI cards with LSI 106x chips.

Obtaining Utilities

The LSI BIOS utility is automatically available in your server's BIOS if a 106x chip is present, either embedded on your server or on a PCI card.

The LSI MegaRAID Storage Manager software should be on your product's Tools and Drivers CD. Alternatively, you can download a CD image from the Sun web site at:

http://sun.com/downloads/

Free registration is required.

On this web page, look for the link labelled "x64 Servers and Workstations". The linked page provides links to all x64-related downloads, organized by product name.

`raidctl` is included in the Solaris OS.

Related Documentation

For all Sun hardware documentation, go to:

`http://www.sun.com/documentation`

For Solaris and other software documentation, go to:

http://docs.sun.com

Third-Party Web Sites

Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused by or in connection with the use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Typographic Conventions

Typeface*	Meaning	Examples
`AaBbCc123`	The names of commands, files, and directories; onscreen computer output	Use `dir` to list all files.
`AaBbCc123`	What you type, when contrasted with onscreen computer output	`> ` **`ipconfig`** `Password:`
AaBbCc123	Book titles, new words or terms, words to be emphasized. Replace command-line variables with real names or values.	Read Chapter 6 in the *User's Guide*. These are called *class* options. You *must* have administrator privileges to do this. To delete a file, type `del` *filename*.
AaBbCc123	Titles of dialog boxes, text in dialog boxes, options, menu items and buttons.	**1. On the** File **menu, click** Extract All**.**

* The settings on your browser might differ from these settings.

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can submit your comments by going to:

`http://www.sun.com/hwdocs/feedback`

Please include the title and part number of your document with your feedback:

Sun LSI 106x RAID User's Guide, 820-4933-12

PART I BIOS RAID Configuration Utility

This part describes how to use the BIOS RAID Configuration utility and has the following chapters:

- "Introduction to Integrated RAID" on page 1
- "Overview of Integrated Mirroring and Integrated Mirroring Enhanced" on page 3
- "Creating IM and IME Volumes" on page 11
- "Overview of Integrated Striping" on page 23
- "Creating Integrated Striping Volumes" on page 27

CHAPTER 1

Introduction to Integrated RAID

This chapter provides an overview of the LSI Integrated RAID solution for LSI SAS controllers. The chapter includes these sections:

- "Introduction" on page 1
- "Integrated RAID Features" on page 2
- "Using this Manual" on page 2

You can use the LSI Integrated RAID solution with the following LSI SAS controllers:

- LSISAS1064/1064E
- LSISAS1068/1068E
- LSISAS1078

Introduction

The components of Integrated RAID are:

- **Integrated Mirroring** (IM), which supports two-disk mirrored arrays and hot-spare disks.
- **Integrated Mirroring Enhanced** (IME), which supports mirrored arrays with three to ten disks, plus hot-spare disks.
- **Integrated Striping** (IS), which supports striped arrays with two to ten disks.

Fusion-MPT™ firmware supports IM, IME, and IS volumes. You can create up to two Integrated RAID storage volumes on the same LSI SAS controller.

Integrated RAID Features

Integrated RAID has the following features:

- Support for up to ten disks per IME or IS volume, with one or two storage volumes per SAS controller. Each controller can support up to 12 volume disks, plus one or two hot-spare disks, for a maximum of 14 disks per controller. (Support for this number of disks requires Integrated RAID firmware v1.20.00 or above.)
- Support for two-disk IM mirrored volumes
- System can boot from an IM, IME, or IS volume.
- Nonvolatile write journaling
- Physical disks are not visible to OS or to application software.
- Functionality is contained in device hardware and firmware.

Using this Manual

Chapters 2 and 3 of this *User's Guide* list IM/IME features and explain how to create IM/IME volumes and optional hot-spare disks.

Chapters 4 and 5 of this *User's Guide* list IS features and explain how to create IS volumes and optional hot-spare disks.

CHAPTER 2

Overview of Integrated Mirroring and Integrated Mirroring Enhanced

This chapter provides an overview of the LSI Integrated Mirroring (IM) and Integrated Mirroring Enhanced (IME) features. It includes these sections:

- "Introduction" on page 3
- "IM and IME Features" on page 4
- "IM/IME Description" on page 5
- "Integrated RAID Firmware" on page 7
- "Fusion-MPT Support" on page 9

Introduction

The LSI Integrated Mirroring (IM) and Integrated Mirroring Enhanced (IME) features provide data protection for the system boot volume to safeguard critical information such as the OS on servers and high-performance workstations. The IM and IME features provide a robust, high-performance, fault-tolerant solution to data storage needs.

The IM and IME features support one or two mirrored volumes per LSI SAS controller, to provide fault-tolerant protection for critical data. The two volumes can have up to twelve disk drives total, plus one or two hot-spare disks.

If a disk in an Integrated Mirroring volume fails, the hot swap capability allows you to restore the volume by simply swapping disks. The firmware then automatically re-mirrors the swapped disk. Additionally, each SAS controller can have one or two global hot-spare disks available to automatically replace a failed disk in the IM or IME storage volumes on the controller. Hot-spares make the IM/IME volume even more fault-tolerant.

Note – You can also configure one IM or IME volume and one Integrated Striping (IS) volume on the same LSI SAS controller.

The IM/IME feature operates independently from the OS, in order to conserve system resources. The BIOS-based configuration utility makes it easy to configure IM and IME volumes.

IM and IME Features

IM and IME support the following features:

1. Configurations of one or two IM or IME volumes on the same LSI SAS controller. IM volumes have two mirrored disks; IME volumes have three to ten mirrored disks. Two volumes can have up to 12 disks total. (Requires Integrated RAID firmware v1.20.00 or above.)
2. One or two global hot-spare disks per controller, to automatically replace failed disks in IM/IME volumes. (Support for two hot-spares requires Integrated RAID firmware v1.20.00 or above.) The hot-spares are in addition to the 12-disk maximum for two volumes per SAS controller.
3. Mirrored volumes run in optimal mode or in degraded mode (if one mirrored disk fails).
4. Hot swap capability
5. Presents a single virtual drive to the OS for each IM/IME volume.
6. Supports both SAS and SATA disks. The two types of disks cannot be combined in the same volume. However, an LSI SAS controller can support one volume with SATA disks and a second volume with SAS disks.
7. Fusion-MPT architecture
8. Easy-to-use BIOS-based configuration utility
9. Error notification: the drivers update an OS-specific event log.
10. SES status LED support
11. Write journaling, which allows automatic synchronization of potentially inconsistent data after unexpected power-down situations.
12. Metadata used to store volume configuration on mirrored disks.
13. Automatic background re-synchronization while host I/O continue.

14. Background media verification ensures that data on IM/IME volumes is always accessible.

IM/IME Description

The LSI Integrated RAID solution supports one or two IM/IME volumes on each LSI SAS controller (or one IM/IME volume and one Integrated Striping volume). Typically, one of these volumes is the primary or boot volume, as shown in FIGURE 2-1. Boot support is available through the firmware of the LSI SAS controller that supports the standard Fusion-MPT interface. The runtime mirroring of the boot disk is transparent to the BIOS, drivers, and OS. Host-based status software monitors the state of the mirrored disks and reports any error conditions. FIGURE 2-1 shows an IM implementation with a second disk as a mirror of the first (primary) disk.

FIGURE 2-1 Typical Integrated Mirroring Implementation

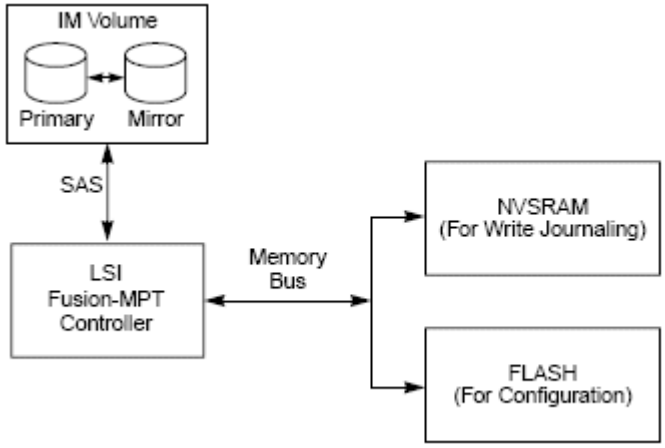

The advantage of an IM/IME volume is that there is always a second, mirrored copy of the data. The disadvantage is that writes take longer because data must be written twice. On the other hand, performance is actually improved during reads.

FIGURE 2-2 shows the logical view and physical view of an IM volume.

FIGURE 2-2 Integrated Mirroring Volume

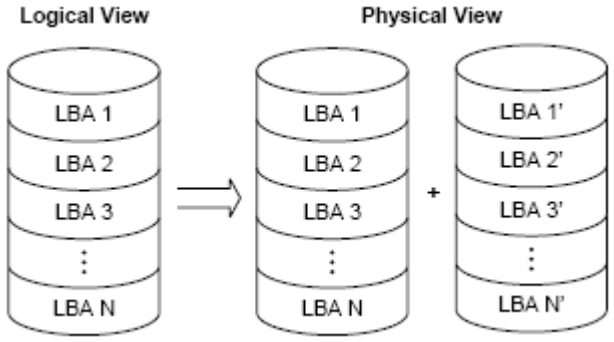

An IME volume can be configured with up to ten mirrored disks. (One or two global hot-spares can be added also.) FIGURE 2-3 shows the logical view and physical view of an Integrated Mirroring Enhanced (IME) volume with three mirrored disks. Each *mirrored stripe* is written to a disk and mirrored to an adjacent disk. This type of configuration is also called RAID 1E.

FIGURE 2-3 Integrated Mirroring Enhanced with Three Disks

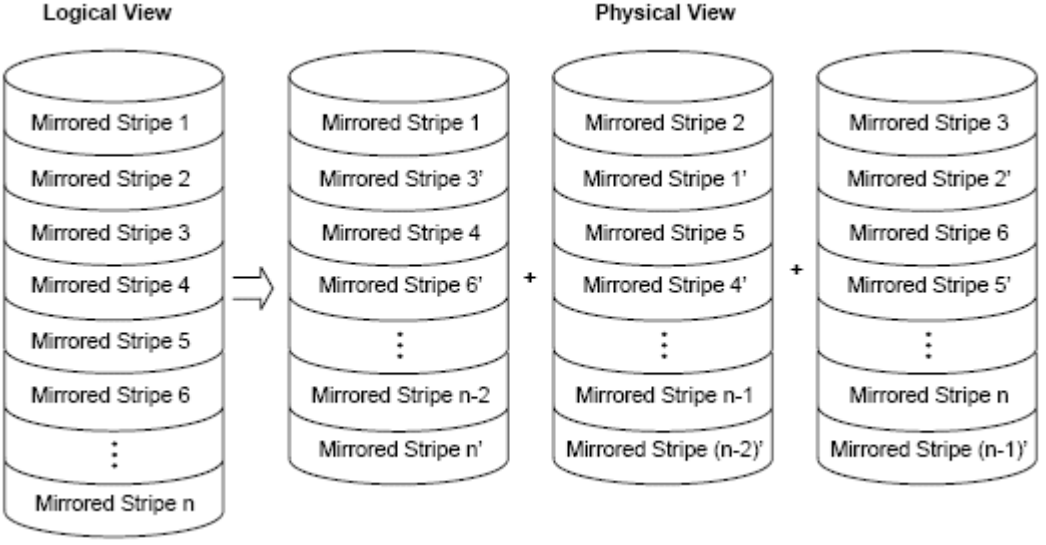

The LSI BIOS-based configuration utility enables you to create IM and IME volumes during initial setup and to reconfigure them in response to hardware failures or changes in the environment.

Integrated RAID Firmware

This section describes features of the LSI Integrated RAID firmware.

Re-synchronization with Concurrent Host I/O Operation

The Integrated RAID firmware allows host I/O to continue on an IM or IME volume while the volume is being re-synchronized in the background. Re-synchronization is attempted after a hot-spare is activated due to a physical device failure, or after a hot swap has occurred to a physical disk in the volume.

Metadata Support

The firmware supports metadata, which describes the IM/IME logical drive configuration stored on each member disk. When the firmware is initialized, each member disk is queried to read the stored metadata in order to verify the configuration. The usable disk space for each member disk is adjusted down when the configuration is created, in order to leave room for this data.

Hot Swapping

The firmware supports *hot swapping*. The hot-swapped disk is automatically re-synchronized in the background, without any host or user intervention. The firmware detects hot swap removal and disk insertion.

Following a hot swap event, the firmware readies the new physical disk by spinning it up and verifying that it has enough capacity for the mirrored volume. The firmware re-synchronizes all hot-swapped disks that have been removed, even if the same disk is re-inserted. In a two-disk mirrored volume, the firmware marks the hot-swapped disk as the secondary disk and marks the other mirrored disk as the primary disk. The firmware re-synchronizes all data from the primary disk onto the new secondary disk.

SMART Support

SMART is a technology that monitors hard disk drives for signs of future disk failure and generates an alert if such signs are detected. The firmware polls each physical disk in the volume at regular intervals. If the firmware detects a SMART ASC/ASCQ code on a physical disk in the IM/IME volume, it processes the SMART data and stores it in nonvolatile memory. The IM/IME volume does not support SMART directly, since it is just a logical representation of the physical disks in the volume.

Hot Spare Disk

One or two disk drives per controller can be configured as *global hot-spare* disks, to protect data on the IM/IME volumes configured on the controller. If the firmware fails one of the mirrored disks, it automatically replaces the failed disk with a hot-spare disk and then re-synchronizes the mirrored data. The firmware is automatically notified when the failed disk has been replaced, and it then designates the failed disk as the new hot-spare.

Media Verification

The firmware supports a background *media verification feature* that runs at regular intervals when the IM/IME volume is in optimal state. If the verification command fails for any reason, the other disk's data for this segment is read and written to the failing disk in an attempt to refresh the data. The current Media Verification Logical Block Address is written to nonvolatile memory occasionally to allow media verification to continue approximately where it left off prior to a power-cycle.

Disk Write Caching

The firmware disables disk write caching by default for IM/IME volumes. This is done to increase data integrity, so that the disk write log stored in NVSRAM is always valid. If disk write caching were enabled (not recommended), the disk write log could be invalid.

Write Journaling

The Integrated RAID firmware requires at least a 32K NVSRAM in order to perform write journaling. Write journaling is used to verify that the disks in the IM/IME volume are synchronized with each other.

Fusion-MPT Support

The BIOS uses the LSI Fusion-MPT interface to communicate to the SAS controller and firmware to enable IM and IME volumes. This includes reading the Fusion-MPT configuration to access the parameters that are used to define behavior between the SAS controller and the devices connected to it. The Fusion-MPT drivers for all supported operating systems implement the Fusion-MPT interface to communicate with the controller and firmware.

CHAPTER 3

Creating IM and IME Volumes

This chapter explains how to create Integrated Mirroring (IM) and Integrated Mirroring Enhanced (IME) volumes using the LSI SAS BIOS Configuration Utility (SAS BIOS CU). The chapter includes these topics:

- "IM/IME Configuration Overview" on page 11
- "Creating IM and IME Volumes" on page 12
- "Creating a Second IM or IME Volume" on page 15
- "Managing Hot Spares" on page 16
- "Other Configuration Tasks" on page 18

IM/IME Configuration Overview

You can use the SAS BIOS CU to create one or two IM/IME volumes on each LSI SAS controller, with one or two optional global hot-spare disks. All disks in an IM/IME volume must be connected to the same LSI SAS controller.

Although you can use disks of different size in IM and IME volumes, the smallest disk in the volume will determine the *logical* size of all disks in the volume. In other words, the excess space of the larger member disk(s) will not be used. For example, if you create an IME volume with two 100 Gbyte disks and two 120 Gbyte disks, only 100 Gbytes of the larger disks will be used for the volume.

Refer to "IM and IME Features" on page 4 for more information about Integrated Mirroring volumes.

Creating IM and IME Volumes

The SAS BIOS CU is part of the Fusion-MPT BIOS. When the BIOS loads during boot and you see the message about the LSI Configuration Utility, press Ctrl-C to start the CU. After you do this, the message changes to:

```
Please wait, invoking SAS Configuration Utility...
```

After a brief pause, the main menu of the SAS BIOS CU appears. On some systems, however, the following message appears next:

```
LSI Configuration Utility will load following initialization!
```

In this case, the SAS BIOS CU will load after the system has completed its POST.

You can configure one or two IM or IME volumes per Fusion-MPT controller. You can also configure one IM/IME and one Integrated Striping (IS) volume on the same controller, up to a maximum of twelve physical disk drives for the two volumes. In addition, you can create one or two hot-spares for the IM/IME array(s).

The following guidelines also apply when creating an IM or IME volume:

- All physical disks in a volume must be either SATA (with extended command set support) or SAS (with SMART support). SAS and SATA disks *cannot* be combined in the same volume. However, you can create one volume with SAS disks and a second volume with SATA disks on the same controller.
- Disks must have 512-byte blocks and must not have removable media.
- An IM volume must have two drives. An IME volume can have three to ten drives. In addition, one or two hot-spares can be created for the IM/IME volume(s).

Note – If a disk in an IM/IME volume fails, it is rebuilt on the global hot-spare if one is available. LSI recommends that you always use hot-spares with IM/IME volumes.

▼ To Create an IM Volume

1. **On the Adapter List screen, use the arrow keys to select an LSI SAS adapter.**
2. **Press Enter to go to the Adapter Properties screen, shown in** FIGURE 3-1**.**

FIGURE 3-1 Adapter Properties Screen

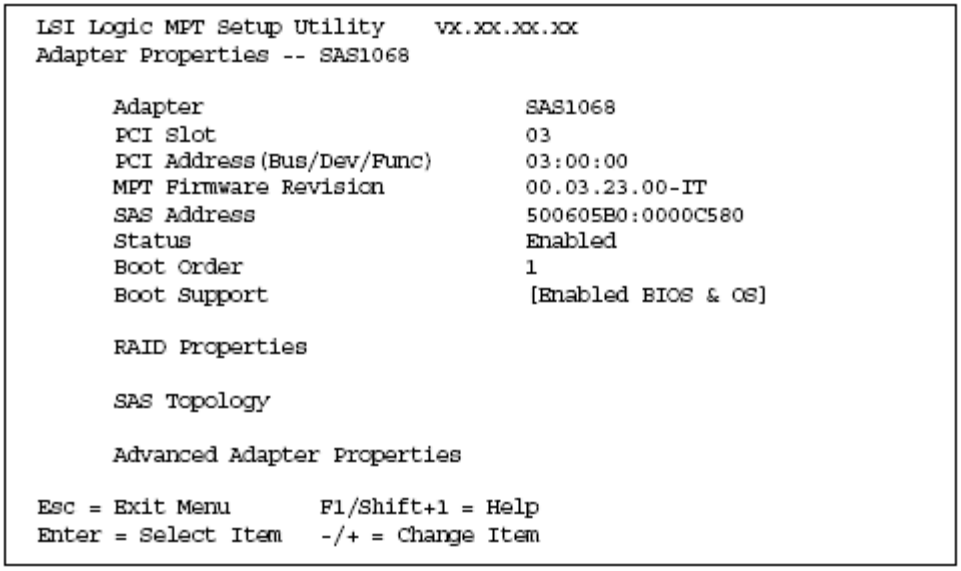

1. **On the Adapter Properties screen, use the arrow keys to select RAID Properties on the screen and press Enter.**
2. **When you are prompted to select a volume type, select Create IM Volume.**

   The Create New Array screen shows a list of disks available to be added to a volume.

FIGURE 3-2 Create New Array Screen

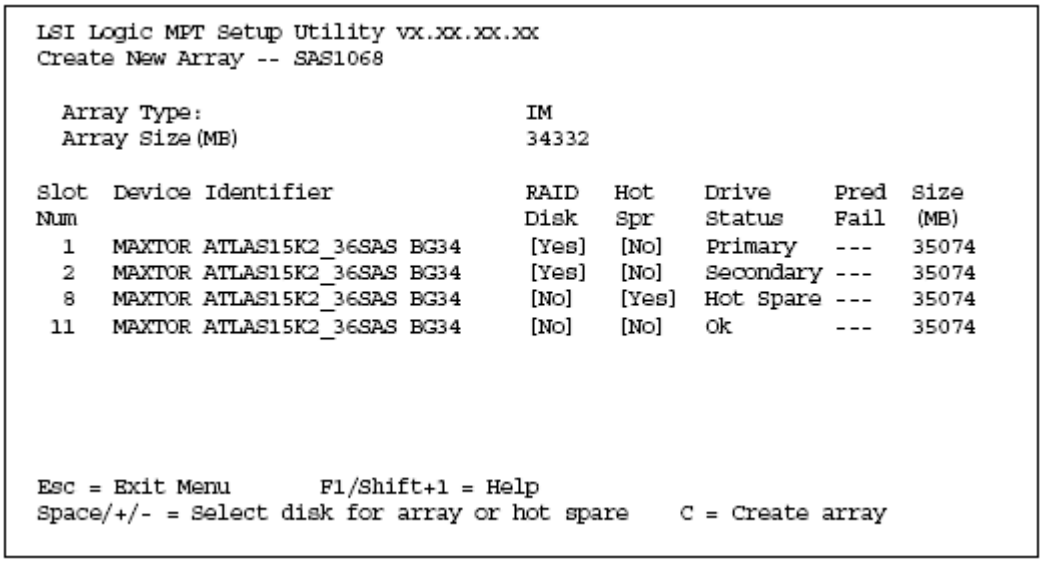
LSI Logic MPT Setup Utility vx.xx.xx.xx
Create New Array -- SAS1068
Array Type: IM
Array Size(MB) 34332
Slot Num
Device Identifier
RAID Disk
Hot Spr
Drive Status
Pred Fail
Size (MB)
1 MAXTOR ATLAS15K2_36SAS BG34 [Yes] [No] Primary --- 35074
2 MAXTOR ATLAS15K2_36SAS BG34 [Yes] [No] Secondary --- 35074
8 MAXTOR ATLAS15K2_36SAS BG34 [No] [Yes] Hot Spare --- 35074
11 MAXTOR ATLAS15K2_36SAS BG34 [No] [No] Ok --- 35074
Esc = Exit Menu F1/Shift+1 = Help
Space/+/- = Select disk for array or hot spare C = Create array

3. **Move the cursor to the** `RAID Disk` **column and select a disk. To add the disk to the volume, change the** `No` **to** `Yes` **by pressing the** `+` **key,** `-` **key, or space bar.**

   When the first disk is added, the SAS BIOS CU prompts you to either keep existing data or overwrite existing data.

4. **Press** `M` **to keep the existing data on the first disk or press** `D` **to overwrite it.**

   If you keep the existing data, this is called a data *migration*. The first disk will be mirrored onto the second disk, so any data you want to keep *must* be on the first disk selected for the volume. Data on the second disk is overwritten. The first disk must have 512 Kbytes available for metadata after the last partition.

   As disks are added the Array Size field changes to reflect the size of the new volume.

5. **[Optional] Add one or two global hot-spares by moving the cursor to the** `Hot Spr` **column and pressing the + key, - key, or space bar.**

   FIGURE 3-2 shows an IM volume configured with one global hot-spare disk.

6. **When the volume has been fully configured, press** `C`**, then select** `Save Changes`**, and then** `Exit This Menu` **to commit the changes.**

   The SAS BIOS CU pauses while the array is being created.

▼ To Create an IME Volume

1. **On the Adapter List screen, use the arrow keys to select an LSI SAS adapter.**
2. **Press Enter to go to the Adapter Properties screen, shown in** FIGURE 3-1**.**
3. **On the Adapter Properties screen, use the arrow keys to select RAID Properties on the screen and press Enter.**
4. **When you are prompted to select a volume type, select Create IME Volume.**

   The Create New Array screen shows a list of disks that can be added to a volume.
5. **Move the cursor to the** `RAID Disk` **column and select a disk. To add the disk to the volume, change the** `No` **to** `Yes` **by pressing the + key, – key, or space bar.**
6. **Repeat this step to select a total of three to ten disks for the volume.**

   All existing data on all the disks you select will be overwritten. As you add disks, the Array Size field changes to reflect the size of the new volume.
7. **[Optional] Add one or two global hot-spares to the volume by moving the cursor to the** `Hot Spr` **column and pressing the + key, – key, or space bar.**
8. **When the volume has been fully configured, press** `C`**, then select** `Save Changes` **and then** `Exit This Menu` **to commit the changes.**

   The SAS BIOS CU pauses while the array is being created.

Creating a Second IM or IME Volume

The LSI SAS controllers allow you to configure two IM or IME volumes per controller. If one volume is already configured, and if there are available disk drives, there are two ways to add a second volume. The first is as follows:

▼ To Create a Second IM or IME Volume Method I

1. **In the configuration utility, select an adapter from the Adapter List, and then select the RAID Properties option.**

   This will display the current volume.
2. **Press** `C` **to create a new volume.**
3. **Continue with** Step 2 **of** "To Create an IM Volume" on page 12 **or** Step 4 **of** "To Create an IME Volume" on page 15 **to create a second volume.**

▼ To Create a Second IM or IME Volume Method II

1. **On the Adapter List screen, use the arrow keys to select an LSI SAS adapter.**
2. **Press Enter to go to the Adapter Properties screen, shown in** FIGURE 3-1**.**
3. **On the Adapter Properties screen, use the arrow keys to select RAID Properties and press Enter.**
4. **Continue with** Step 2 **of** "To Create an IM Volume" on page 12 **or** Step 4 **of** "To Create an IME Volume" on page 15 **to create a second volume.**

Managing Hot Spares

You can create one or two global hot-spare disks to protect the IM or IME volumes on a SAS controller. Usually, you create global hot-spares at the same time you create the IM/IME volume. Follow these steps to add global hot-spare disks later:

▼ To Create Global Hot-Spare Disks

1. **On the View Array screen, select Manage Array.**
2. **Select Manage Hot Spare on the Manage Array screen, shown in** FIGURE 3-3**.**

FIGURE 3-3 Manage Array Screen

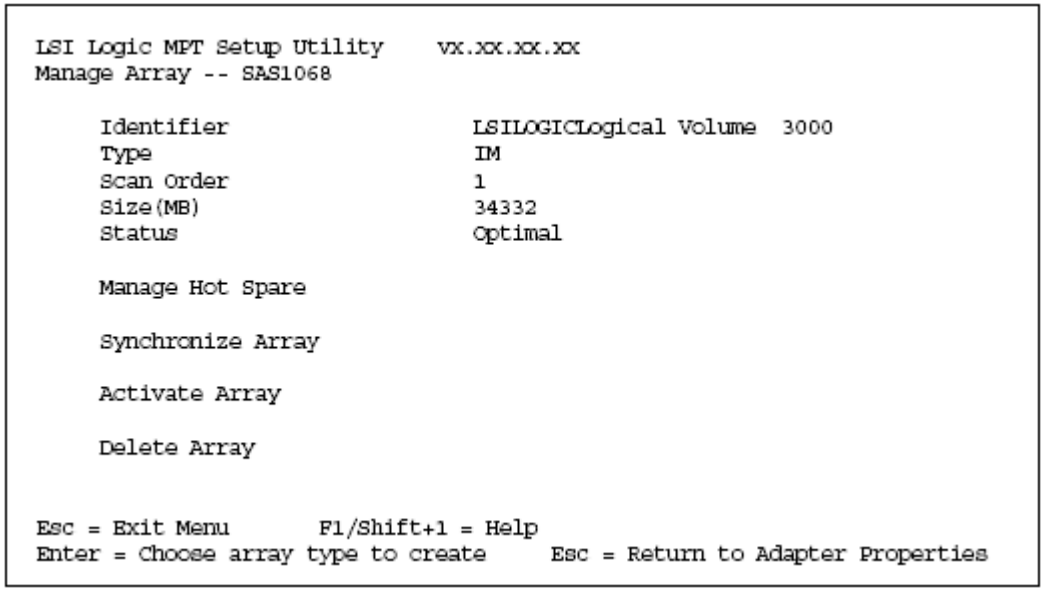
LSI Logic MPT Setup Utility vx.xx.xx.xx
Manage Array -- SAS1068
Identifier LSILOGICLogical Volume 3000
Type IM
Scan Order 1
Size(MB) 34332
Status Optimal
Manage Hot Spare
Synchronize Array
Activate Array
Delete Array
Esc = Exit Menu F1/Shift+1 = Help
Enter = Choose array type to create Esc = Return to Adapter Properties

3. **Select a disk from the list by pressing the + key, – key, or spacebar.**
4. **After you select the global hot-spare disk, press** C**.**

   An error message appears if the selected disk is not at least as large as the smallest disk used in the IM/IME volume(s). The global hot-spare disk must have 512-byte blocks, it cannot have removable media, and the disk type must be either SATA with extended command set support or SAS with SMART support.

   If SATA disks are used for the IM/IME volume(s), the hot-spare disk must also be a SATA disk. If SAS disks are used, the hot-spare disk must also be a SAS disk. An error message appears if the selected disk is not the same type as the disks used in the IM/IME volumes.
5. **[Optional] Select a second hot-spare disk.**
6. **Select Save changes then exit this menu to commit the changes.**

   The configuration utility pauses while the global hot-spares are being added.

Follow these steps to delete a global hot-spare:

1. **Select Manage Hot Spare on the Manage Array screen.**
2. **Select Delete Hot Spare and then press** C**.**

   If there are two hot-spares, select which one you want to delete.

3. **Select Save changes then exit this menu to commit the changes.**

The configuration utility pauses while the global hot-spare is being removed.

Other Configuration Tasks

This section explains how to perform other configuration and maintenance tasks for IM and IME volumes:

- "To View Volume Properties" on page 18
- "Synchronizing an Array" on page 19
- "Activating an Array" on page 19
- "Deleting an Array" on page 20
- "Locating a Disk Drive or Multiple Disk Drives in a Volume" on page 20
- "Selecting a Boot Disk" on page 21

▼ To View Volume Properties

1. **In the SAS BIOS CU, select an adapter from the Adapter List. Select the RAID Properties option.**

   The properties of the current volume are displayed. If global hot-spares are defined, they are also listed.

Note – If you create one volume using SAS disks, another volume using SATA disks, and global hot-spare disks, the hot-spare disks will only appear when you view the volume that has the same type of disks as the hot-spare disks.

---

2. **If two volumes are configured, press Alt+N to view the other array.**
3. **To manage the current array, select the Manage Array item and press Enter.**

Synchronizing an Array

The Synchronize Array command forces the firmware to re-synchronize the data on the mirrored disks in the array. It is seldom necessary to use this command, because the firmware automatically keeps the mirrored data synchronized during normal operation. When you use this command, one disk of the array is placed in the *Degraded* state until the data on the mirrored disks has been re-synchronized.

▼ To Synchronize an Array

1. **Select Synchronize Array on the Manage Array screen.**
2. **Press `Y` to start the synchronization, or `N` to cancel it.**

Activating an Array

An array can become inactive if, for example, it is removed from one controller or computer and moved to another one. The *Activate Array* option allows you to reactivate an inactive array that has been added to a system. This option is only available when the selected array is currently inactive.

▼ To Activate an Array

1. **Select Activate Array on the Manage Array screen.**
2. **Press `Y` to proceed with the activation, or press `N` to abandon it.**

   After a pause, the array will become active.

Note – If there is a global hot-spare disk on the controller to which you have moved the array, the BIOS checks when you activate the array to determine if the hot-spare is compatible with the new array. An error message appears if the disks in the activated array are larger than the hot-spare disk or if the disks in the activated array are not the same type as the hot-spare disk (SATA versus SAS).

Deleting an Array

Caution – Before deleting an array, be sure to back up all data on the array that you want to keep.

▼ To Delete an Array

1. **Select Delete Array on the Manage Array screen.**
2. **Press** `Y` **to delete the array.**

   After a pause, the array is deleted. If there is another remaining array and one or two hot-spare disks, the BIOS checks the hot-spare disks to determine if they are compatible with the remaining array. If they are not compatible (too small or wrong disk type) the firmware deletes them also.

Note – After a volume has been deleted, it cannot be recovered. When an IM volume is deleted, the data is preserved on the primary disk. When an IME volume is deleted, the master boot records of all disks are deleted.

Locating a Disk Drive or Multiple Disk Drives in a Volume

You can use the SAS BIOS CU to locate and identify a specific physical disk drive by blinking the drive's LED. You can also use the SAS BIOS CU to blink the LEDs of all the disk drives in a RAID volume. There are several ways to do this:

1. When you are creating an IM or IME volume, and a disk drive is set to Yes as part of the volume, the LED on the disk drive is blinking. The LED is turned off when you have finished creating the volume.
2. You can locate individual disk drives from the SAS Topology screen. To do this, move the cursor to the name of the disk in the Device Identifier column and press Enter. The LED on the disk blinks until the next key is pressed.
3. You can locate all the disk drives in a volume by selecting the volume on the SAS Topology screen. The LEDs blink on all disk drives in the volume.

Note – The LEDs on the disk drives will blink as described above if the firmware is correctly configured and the drives or the disk enclosure supports disk location.

Selecting a Boot Disk

You can select a boot disk in the SAS Topology screen. This disk is then moved to scan ID 0 on the next boot, and remains at this position. This makes it easier to set BIOS boot device options and to keep the boot device constant during device additions and removals. There can be only one boot disk.

▼ To Select a Boot Disk

1. **In the SAS BIOS CU, select an adapter from the Adapter List.**
2. **Select the SAS Topology option.**

   The current topology appears. If the selection of a boot device is supported, the bottom of the screen lists the Alt+B option. This is the key for toggling the boot device. If a device is currently configured as the boot device, the Device Info column on the SAS Topology screen will show the word "Boot."
3. **To select a boot disk, move the cursor to the disk and press Alt+B.**
4. **To remove the boot designator, move the cursor down to the current boot disk and press Alt+B. This controller will no longer have a disk designated as boot.**
5. **To change the boot disk, move the cursor to the new boot disk and press Alt+B. The boot designator will move to this disk.**

Note – The firmware must be configured correctly in order for the Alt+B feature to work.

CHAPTER 4

Overview of Integrated Striping

This chapter provides an overview of the LSI Integrated Striping (IS) feature. It includes these sections:

- "Introduction" on page 23
- "IS Features" on page 24
- "IS Description" on page 24
- "Integrated Striping Firmware" on page 26
- "Fusion-MPT Support" on page 26

Introduction

The LSI Integrated Striping (IS) feature is useful for applications that require the faster performance and increased storage capacity of striping. The low-cost IS feature has many of the advantages of a more expensive RAID striping solution. A single IS logical drive may be configured as the boot disk or as a data disk.

The IS feature is implemented with controller firmware that supports the Fusion-MPT Interface. IS provides better performance and more capacity than individual disks, without burdening the host CPU. The firmware splits host I/Os over multiple disks and presents the disks as a single logical drive. In general, striping is transparent to the BIOS, the drivers, and the operating system.

The SAS BIOS CU is used to configure IS volumes, which can consist of two to ten disks.

IS Features

IS includes the following features:

- Support for volumes with two to ten disks
- Support for two IS volumes (or one IS volume and one IM/IME volume) on a controller, with up to 12 disks total (Requires Integrated RAID firmware v1.20.00 or above.)

Note – All physical disks in a volume must be connected to the same SAS controller.

- Presents a single virtual drive to the OS for each configured volume
- Support for both SAS and SATA drives, although the two types of drives cannot be combined in one volume
- Use of metadata to store volume configuration on disks
- SES status LED support for drives used in IS volumes

IS Description

The IS feature writes data across multiple disks instead of onto one disk. This is accomplished by partitioning each disk's storage space into 64 Kbyte stripes. These stripes are interleaved round-robin, so that the combined storage space is composed alternately of stripes from each disk.

For example, as shown in FIGURE 4-1, segment 1 is written to disk 1, segment 2 is written to disk 2, segment 3 is written to disk 3, and so on. When the system reaches the end of the disk list, it continues writing data at the next available segment of disk 1.

FIGURE 4-1 Integrated Striping Example

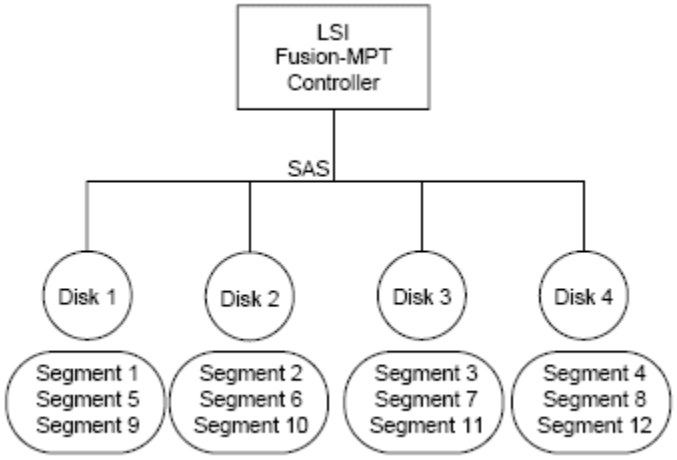

FIGURE 4-2 shows a logical view and a physical view of Integrated Striping configuration.

FIGURE 4-2 Integrated Striping - Logical and Physical Views

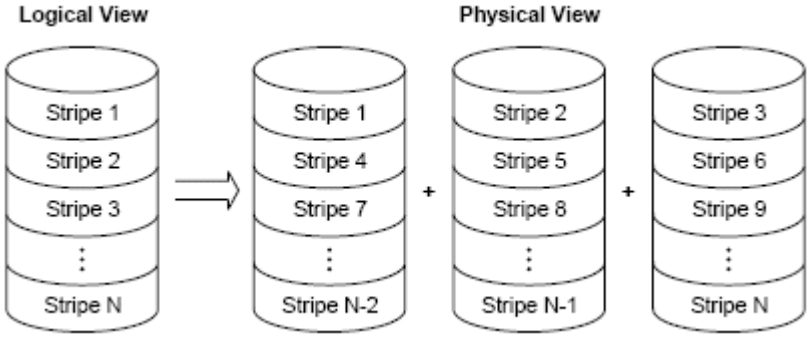

The primary advantage of IS is speed, because it transfers data to or from multiple disks at once. However, there is no data redundancy; therefore, if one disk fails, that data is lost.

Integrated Striping Firmware

This section describes features of the LSI Integrated RAID firmware.

Metadata Support

The firmware supports metadata, which describes the IS logical drive configuration stored on each member disk. When the firmware is initialized, each member disk is queried to read the stored metadata to verify the configuration. The usable disk space for each IS member disk is adjusted down when the configuration is created, in order to leave room for this data.

SMART Support

SMART is a technology that monitors disk drives for signs of future disk failure and generates an alert if such signs are detected. The firmware polls each physical disk in the volume at regular intervals. If the firmware detects a SMART ASC/ASCQ code on a physical disk in the IS volume, it processes the SMART data and stores it in nonvolatile memory. The IS volume does not support SMART directly, since it is just a logical representation of the physical disks in the volume.

Disk Write Caching

Disk write caching is enabled by default on all IS volumes.

Fusion-MPT Support

The BIOS uses the LSI Fusion-MPT interface to communicate to the SAS controller and firmware to enable IS. This includes reading the Fusion-MPT configuration to gain access to the parameters that are used to define behavior between the SAS controller and the devices connected to it. The Fusion-MPT drivers for all supported operating systems implement the Fusion-MPT interface to communicate with the controller and firmware.

CHAPTER 5

Creating Integrated Striping Volumes

This chapter explains how to create Integrated Striping (IS) volumes using the LSI SAS BIOS Configuration Utility (SAS BIOS CU). The chapter includes these topics:

- "IS Configuration Overview" on page 27
- "Creating IS Volumes" on page 28
- "Creating a Second IS Volume" on page 30
- "Other Configuration Tasks" on page 31

IS Configuration Overview

You can use the SAS BIOS CU to create one or two IS volumes, with up to twelve drives total, on an LSI SAS controller. Each volume can have from two to ten drives. Disks in an IS volume must be connected to the same LSI SAS controller, and the controller must be in the BIOS boot order.

Although you can use disks of different size in IS volumes, the smallest disk determines the "logical" size of each disk in the volume. In other words, the excess space of the larger member disk(s) is not used. Usable disk space for each disk in an IS volume is adjusted down to leave room for metadata. Usable disk space may be further reduced to maximize the ability to interchange disks in the same size classification. The supported stripe size is 64 kilobytes.

Refer to "IS Features" on page 24 for more information about Integrated Striping volumes.

Creating IS Volumes

The SAS BIOS CU is part of the Fusion-MPT BIOS. When the BIOS loads during boot and you see the message about the LSI Configuration Utility, press Ctrl-C to start it. After you do this, the message changes to:

```
Please wait, invoking SAS Configuration Utility...
```

After a brief pause, the main menu of the SAS BIOS CU appears. On some systems, however, the following message appears next:

```
LSI Logic Configuration Utility will load following
initialization!
```

In this case, the SAS BIOS CU will load after the system has completed its power-on self test.

Follow the steps below to configure an Integrated Striping (IS) volume with the SAS BIOS CU. The procedure assumes that the required controller(s) and disks are already installed in the computer system. You can configure an IM/IME volume and an IS volume on the same SAS controller.

▼ To Create IS Volumes

1. **On the Adapter List screen of the SAS BIOS CU, use the arrow keys to select a SAS adapter.**
2. **Press Enter to go to the Adapter Properties screen, shown in** FIGURE 5-1**.**

FIGURE 5-1 Adapter Properties Screen

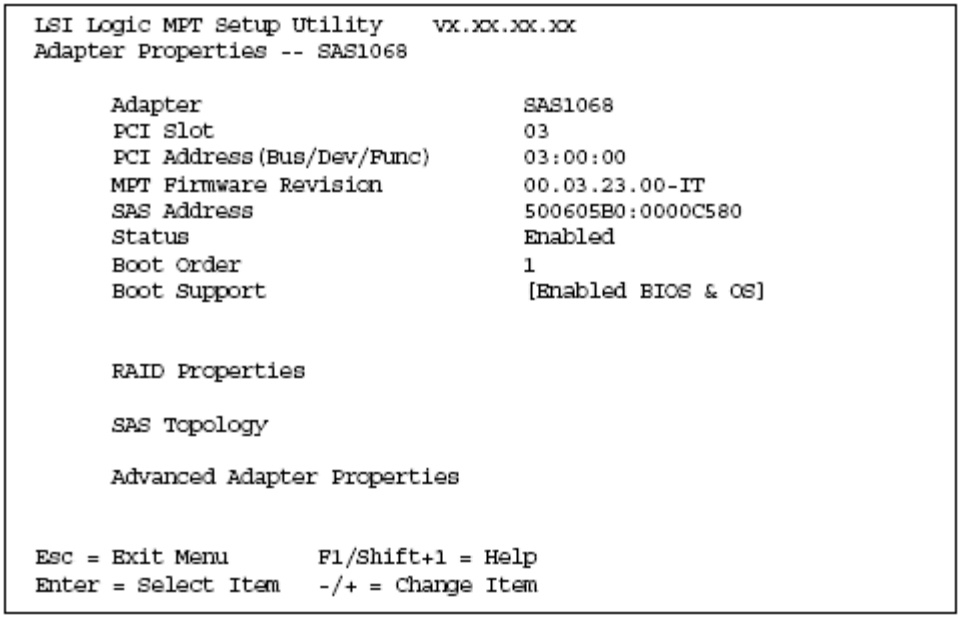

3. **On the Adapter Properties screen, use the arrow keys to select RAID Properties and press Enter.**
4. **When you are prompted to select a volume type, select** `Create IS Volume`**.**
5. **The Create New Array screen shows a list of disks that can be added to a volume.**

FIGURE 5-2 shows the Create New Array screen with an IS volume configured with two drives.

FIGURE 5-2 Create New Array Screen

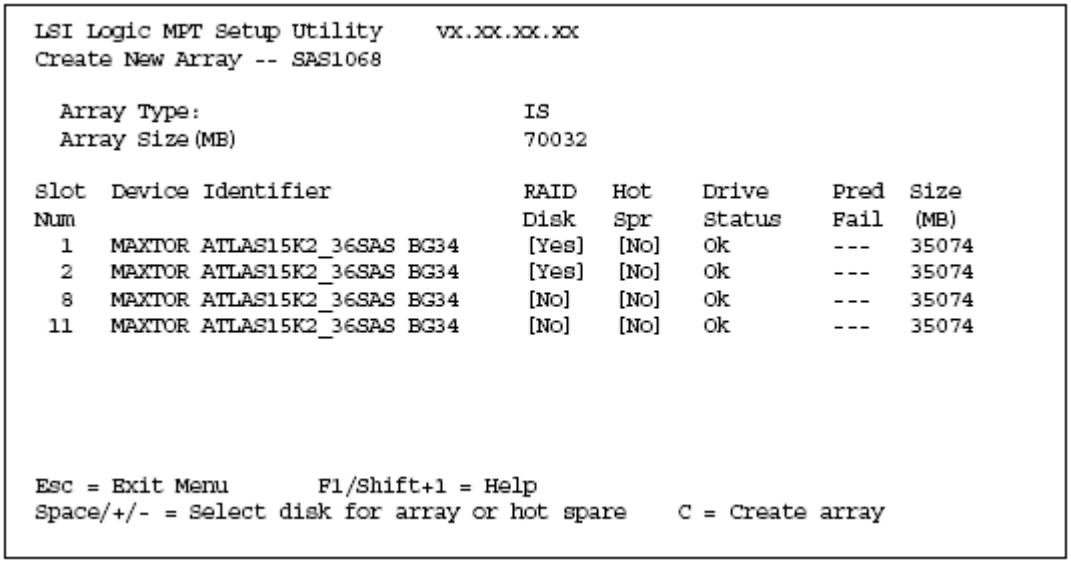
LSI Logic MPT Setup Utility vx.xx.xx.xx
Create New Array -- SAS1068
Array Type: IS
Array Size(MB) 70032
Slot Num
Device Identifier
RAID Disk
Hot Spr
Drive Status
Pred Fail
Size (MB)
1 MAXTOR ATLAS15K2_36SAS BG34 [Yes] [No] Ok --- 35074
2 MAXTOR ATLAS15K2_36SAS BG34 [Yes] [No] Ok --- 35074
8 MAXTOR ATLAS15K2_36SAS BG34 [No] [No] Ok --- 35074
11 MAXTOR ATLAS15K2_36SAS BG34 [No] [No] Ok --- 35074
Esc = Exit Menu F1/Shift+1 = Help
Space/+/- = Select disk for array or hot spare C = Create array

6. **Move the cursor to the** `RAID Disk` **column. To add a disk to the volume, change the** `No` **to** `Yes` **by pressing the + key, – key, or space bar. As disks are added, the** `Array Size` **field changes to reflect the size of the new volume.**

   There are several limitations when creating an IS (RAID 0) volume:

   - All disks must be either SATA (with extended command set support) or SAS (with SMART support).
   - Disks must have 512-byte blocks and must not have removable media.
   - There must be at least two and no more than ten drives in a valid IS volume. Hot-spare drives are not allowed.

4. When you have added the desired number of disks to the array, press `C`, then select `Save Changes, and then Exit This Menu` to commit the changes. The configuration utility pauses while the array is being created.

Creating a Second IS Volume

The LSI SAS controllers allow you to configure two IS volumes, or an IS volume and an IM/IME volume. If one volume is already configured, and if there are available disk drives, there are two ways to add a second volume.

▼ To Create a Second IS Volume, Method I

1. **In the configuration utility, select an adapter from the Adapter List, and then select the RAID Properties option.**

   This will display the current volume.

2. **Press C to create a new volume.**

3. **Continue with** Step 4 **of** "Creating IS Volumes" on page 28 **to create a second IS volume.**

▼ To Create a Second IS Volume, Method II

1. **On the Adapter List screen, use the arrow keys to select an LSI SAS adapter.**

2. **Press Enter to go to the Adapter Properties screen, shown in** FIGURE 5-1**.**

3. **On the Adapter Properties screen, use the arrow keys to select RAID Properties and press Enter.**

4. **Continue with** Step 4 **of** "Creating IS Volumes" on page 28 **to create a second IS volume.**

Other Configuration Tasks

This section explains how to perform other configuration and maintenance tasks for IS volumes.

- "To View IS Volume Properties" on page 31
- "Activating an Array" on page 32
- "To Delete an Array" on page 32
- "Locating a Disk Drive or Multiple Disk Drives in a Volume" on page 33
- "Selecting a Boot Disk" on page 33

▼ To View IS Volume Properties

1. **In the configuration utility, select an adapter from the Adapter List. Select the RAID Properties option.**

   The properties of the current volume appears.

2. **If more than one volume is configured, press Alt+N to view the next array.**
3. **To manage the current array, press Enter when the Manage Array item is selected.**

Activating an Array

An array can become inactive if, for example, it is removed from one controller or computer and moved to another one. The "Activate Array" option allows you to reactivate an inactive array that has been added to a system. This option is only available when the selected array is currently inactive.

▼ To Activate a Selected Array

1. **Select Activate Array on the Manage Array screen.**
2. **Press `Y` to proceed with the activation, or press `N` to abandon it.**

   After a pause, the array will become active.

▼ To Delete an Array

Caution – Before deleting an array, be sure to back up all data on the array that you want to keep.

1. **Select Delete Array on the Manage Array screen.**
2. **Press `Y` to delete the array, or press `N` to abandon the deletion.**

   After a pause, the firmware deletes the array.

Note – Once a volume has been deleted, it cannot be recovered. The master boot records of all disks are deleted.

Locating a Disk Drive or Multiple Disk Drives in a Volume

You can use the SAS BIOS CU to locate and identify a specific physical disk drive by flashing the drive's LED. You can also use the SAS BIOS CU to flash the LEDs of all the disk drives in a RAID volume. There are several ways to do this:

1. When you are creating an IS volume, and a disk drive is set to Yes as part of the volume, the LED on the disk drive is flashing. The LED is turned off when you have finished creating the volume.
2. You can locate individual disk drives from the SAS Topology screen. To do this, move the cursor to the name of the disk in the Device Identifier column and press Enter. The LED on the disk flashes until the next key is pressed.
3. You can locate all the disk drives in a volume by selecting the volume on the SAS Topology screen. The LEDs flash on all disk drives in the volume.

Note – The LEDs on the disk drives will flash as described above if the firmware is correctly configured and the drives or the disk enclosure supports disk location.

Selecting a Boot Disk

You can select a boot disk in the SAS Topology screen. This disk is then moved to scan ID 0 on the next boot, and remains at this position. This makes it easier to set BIOS boot device options and to keep the boot device constant during device additions and removals. There can be only one boot disk.

▼ To Select a Boot Disk

1. **In the SAS BIOS CU, select an adapter from the Adapter List.**
2. **Select the SAS Topology option.**

   The current topology is displayed. If the selection of a boot device is supported, the bottom of the screen lists the Alt+B option. This is the key for toggling the boot device. If a device is currently configured as the boot device, the Device Info column on the SAS Topology screen will show the word "Boot."
3. **To select a boot disk, move the cursor to the disk and press Alt+B.**
4. **To remove the boot designator, move the cursor down to the current boot disk and press Alt+B. This controller will no longer have a disk designated as boot.**

5. **To change the boot disk, move the cursor to the new boot disk and press Alt+B. The boot designator will move to this disk.**

Note – The firmware must be configured correctly in order for the Alt+B feature to work.

PART II LSI `cfggen` Utility

This part describes how to use the LSI `cfggen` utility:

- "The LSI `cfggen` Utility" on page 6-37

CHAPTER 6

The LSI `cfggen` Utility

The LSI `cfggen` utility is a configuration utility used to create Integrated Mirroring (IM) volumes.

The chapter has the following sections:

- "Installing the `cfggen` Utility" on page 37
- "Overview of `cfggen`" on page 38
- "`cfggen` Syntax" on page 38
- "Common Command-Line Parameters" on page 39
- "Supported Commands" on page 40
- "Monitoring and Managing RAID Arrays" on page 53

Installing the `cfggen` Utility

You can install the utility from the Tools and Drivers CD, if available, or from your product Tools and Drivers CD image, downloadable from the product web page.

The `cfggen` utility resides in `/windows/w2k3/tools/lsi_cfggen`. The directory includes 32-bit and 64-bit executables as well as LSI documentation.

Note – This chapter describes the utility *as implemented on Sun's x64 servers*. The LSI documentation in the Tools and Drivers CD describes the utility in general.

Overview of `cfggen`

The `cfggen` utility is a configuration utility used to create Integrated Mirroring (IM) volumes. A command-line utility, it runs in the Windows Preinstallation Environment (WinPE) and on DOS. The utility is a minimally interactive program that can be executed from a command-line prompt or a shell script. The result of running this utility is communicated through the program status value that is returned when the program exits. You use the utility to create IM storage configurations on both SCSI controllers and SAS controllers.

The utility runs on WinPE and is statically compiled with the LSI MptLib Library (`MptLib.lib`). The WinPE environment must have the appropriate LSI Logic MPT Windows driver (`ScsiPort` or `StorPort`) installed and loaded in order to recognize and communicate with the I/O controller. The utility does not recognize an LSI53C1030 or LSI53C1020 controller unless there is at least one device attached to the controller.

Caution – Do not run `cfggen` in a command-line window from within Windows.

`cfggen` Syntax

`cfggen` uses a command line interface with the following format:

```
cfggen controller-number command-parameters
```

Note the following:

- Information is passed between the user environment and `cfggen` through the command line, the standard output, and standard error interfaces, and the program return value.
- You can redirect the output streams as permitted by the operating environment.
- The program return value is returned when the program exits.
- A value of 0 is returned if the command is successful. Otherwise, a value of 1 is returned.

`cfggen` Command Conventions

The following conventions are used in the command descriptions:

- `cfggen` is not case sensitive. Commands and parameters can be typed in uppercase, lowercase, or a mixture of the two.
- Text in italics must be entered exactly as shown on the command line
- Text surrounded by [ ] may be replaced by an optional parameter
- Parameters surrounded by {} must be entered one or more times, as is appropriate for the command being executed
- The command- line definition characters <>, [ ], and {} must not be entered on the command line

Common Command-Line Parameters

The following `cfggen` command line parameters are common to more than one command:

- *controller-number*

The unique number of a PCI function found in the system, starting with controller number 0. Therefore, the *controller-number* is used to address a particular SCSI bus in the system. For example, cfggen assigns two controller numbers to an LSI53C1030 dual SCSI bus chip. It assigns one controller number to an LSI53C1020 single SCSI bus chip. For the LSI Logic SAS1064/1064E and SAS1068/1068E controllers, the controller number corresponds to a single SAS controller.

For example, with `cfggen` in a system containing two SAS1068 controllers, controller number 0 references the first controller and controller number 1 references the other controller.

Valid controller number values are 0 to 255 (decimal).

- *SCSI-ID*

The SCSI bus address of a peripheral device attached to an LSI Logic controller. The maximum value of SCSI ID depends on the type of I/O controller and the maximum number of devices supported by the OS for this controller.

Valid SCSI ID values are:

0-15 (decimal) per SCSI bus for LSI53C1020/1030 controllers

0-127 (decimal) per controller for SAS1064/1064E and SAS1068/1068E controllers

Note – With PBSRAM, the SAS1068/1068E controllers can support more than 128 devices.

- *enclosure:bay*

The enclosure (encl) and bay/slot of a peripheral device attached to the bus. The argument must use a colon (:) as a separator and must follow the enclosure:bay format. Only devices connected to LSI SAS controllers can be addressed using enclosure:bay and hence this option is not supported on LSI53C1020/1030 controllers.

Valid numbers are:

enclosure: A 16-bit EnclosureHandle value set by the I/O controller. A value of 0 is invalid.

bay/slot: A 16-bit Slot value set by the I/O controller.

The enclosure and slot numbers of a drive can be obtained from the `display` command

Supported Commands

The following commands are currently supported by `cfggen`:

- "`auto` Command" on page 40
- "`create` Command" on page 42
- "`display` Command" on page 44
- "`delete` Command" on page 44
- "`hotspare` command" on page 48
- "`list` command" on page 49
- "`rebuild` command" on page 50
- "`status` Command" on page 51

`auto` Command

The AUTO command automatically creates an IM, IME, or IS volume on an LSI1064/1064E or LSI1068/1068E controller. The volume is created with the maximum number of disks available for use in the specified volume type. The main difference from the CREATE command is that with the AUTO command you do not specify SCSI ID values for disks to use in the volume. CFGGEN automatically

creates the volume with the first usable disks it finds. Firmware and hardware limitations for the family of controllers limit the number of configurations that are possible.

When a disk drive is added to an IM, IME, or IS volume, its entire storage capacity may or may not be used, depending on drive capacity and volume capacity. For example, if you add a 36 Gbyte disk drive to a volume that only uses 9 Gbytes of capacity on each disk drive, the remaining 27 Gbytes of capacity on the disk drive are unusable. When AUTO creates an IM volume, the first disk found is assigned as the primary disk drive. If the controller is allowed to resync the disk drives, the data on the primary disk drive will be available by accessing the newly created volume.

CFGGEN follows these rules when creating IM, IME, and IS volumes and hot spare disks with the AUTO command:

- All disks that are part of a volume or a hot spares for a volume must be connected to the same controller.
- IM, IME, and IS volumes are supported.
- Only two volumes per controller can be created.
- SAS and SATA drives cannot be mixed in a volume. With the AUTO command, all drives used must be the same type as the first available disk found.
- The total number of disks in a volume, including hot spare disks, cannot exceed eight for LSI1064/1064E and LSI1068/1068E controllers, and the total number of disks combined for two volumes cannot exceed ten. An IM volume must have exactly two disks.
- An IME volume can have three to six disks for an LSI SCSI controller, and three to eight disks for an LSI SAS controller as long as rules 4 and 5 are not violated.

Example

cfggen *controller-number* **auto** *volume-type size size* `[qsync] [noprompt]`

Parameters

controller-number Number of the SAS controller targeted by this command.

Note – Specifying SAS instead of controller # will configure all SAS only controllers in the system and is supported only for AUTO command in Linux version.

volume-type	Volume type for the volume to be created. Valid values are IM, IME and IS
size	Size of the RAID volume in Mbytes, or MAX for the maximum size available
qsync	If this optional parameter is specified, a quick synchronization of the new volume will be performed. If the volume type is IME or IS, a quick synchronization is always performed even if this option is not specified. A quick synchronization means that the first 32 Kbytes of the drives in the volume are cleared to 0.
noprompt	Suppresses display of warnings and prompts

Program Return Value

0x00 SUCCESS command completed successfully
0x01 FAILURE bad command line arguments or operational failure
0x02 ADAPTER_NOT_FOUND adapter specified cannot be found

`create` Command

The `create` command creates IM, IME (Integrated Mirroring Enhanced), and IS (Integrated Striping) volumes on the LSI53C1020/1030 and SAS1064/1064E and SAS1068/1068E controllers. The firmware and hardware limitations for these controllers determine the number of configurations that can be created. When a disk drive is added to an IM, IME, or IS volume, its entire storage capacity can be used, depending on drive capacity and volume capacity. Any unused capacity is not accessible for other purposes.

For example, if you add a 36 Gbyte disk drive to a volume that only uses 9 Gbytes of capacity on each disk drive, the remaining 27 Gbytes of capacity on the disk drive is unusable. The disk identified by the first SCSI ID on the command line is assigned as the primary disk drive when an IM volume is created. If the controller is allowed to resynchronize the disk drives, the data on the primary disk drive will be available when you access the newly created volume.

The following rules must be observed when creating IM, IME, and IS volumes and hot spare disks:

1. All disks that are part of a volume, including hot spares for that volume, must be on the same SAS controller or on the same SCSI bus (for SCSI controllers).

2. IM, IME, and IS volumes are supported.
3. A maximum of two IM, IME, or IS volumes per controller can be created.
4. The total number of disks in a volume, including hot-spare disks, cannot exceed six for LSI53C1020/1030 controllers.
5. The total number of disks in a volume, including hot-spare disks, cannot exceed eight for SAS1064/1064E and SAS1068/1068E controllers, and the total number of disks combined for two volumes cannot exceed ten. Ten disks is a theoretical upper limit for the firmware; the SAS controller can actually support a fewer number of disks.
6. An IM volume must have exactly two disks.
7. An IME volume can have a minimum of three disks and a maximum of six disks (for LSI53C1020/1030 controllers) or eight disks (for SAS controllers), as long as rules 4 and 5 are not violated.

Example

cfggen *controller-number* **create** *volume-type size* {*SCSI-ID*} [qsync] [noprompt]

cfggen *controller-number* **create** *volume-type size* {*encl:bay*} [qsync] [noprompt]

controller-number	Number of the SCSI bus or SAS controller targeted by this command
volume-type	Volume type for the new volume to be created. Valid values are IM or IME or IS.
size	Size of the RAID volume in Mbytes, or MAX for the maximum size available
SCSI-ID	SCSI ID of a hard drive to be included in the RAID volume
encl:bay	The enclosure:bay value for the disk drive to be included in the RAID volume. These values can be obtained from the output of the DISPLAY command.
qsync	If this optional parameter is specified, a quick synchronization of new volume will be performed. If the volume type is IME or IS, a quick synchronization is always performed even if qsync is not specified. A quick synchronization means that the first 32 Kbytes of the drives in the volume are cleared to 0
noprompt	Suppresses display of warnings and prompts

Program Return Value

0x00 SUCCESS command completed successfully
0x01 FAILURE bad command line arguments or operational failure
0x02 ADAPTER_NOT_FOUND adapter specified cannot be found

`delete` Command

The `delete` command deletes all IM, IME, and IS volumes and hot spare drives. No other controller configuration parameters are changed.

Example

`cfggen` *controller-number* **`delete`** `[noprompt]`

controller-number	Number of the SCSI bus or SAS controller targeted by this command
`noprompt`	Suppresses display of warnings and prompts

Program Return Value

0x00 SUCCESS command completed successfully
0x01 FAILURE bad command line arguments or operational failure
0x02 ADAPTER_NOT_FOUND adapter specified cannot be found

`display` Command

The `display` command displays configuration information for the supported LSI controllers. The information includes controller type, firmware version, BIOS version (version executed), volume information, and physical drive information. An example of the information that will be output by this command is provided below.

Note – 1 Mbyte = 1,048,576 bytes. All sizes displayed in Mbytes are rounded down to the nearest Mbyte.

Example

cfggen *controller-number* **display** [*filename*]

controller-number	Number of the SCSI bus or SAS controller targeted by this command
filename	Optional valid filename to store output of command to a file

Program Return Value

0x00 SUCCESS command completed successfully
0x01 FAILURE bad command line arguments or operational failure
0x02 ADAPTER_NOT_FOUND adapter specified cannot be found

Sample Output

The following example shows the output of the `display` command with an IM configuration on a SAS1068 controller.

Note – The format and content of the display command output might vary depending on the version being used.

```
Read configuration has been initiated for controller 0
------------------------------------------------------------------------
Controller information
------------------------------------------------------------------------
   Controller type : SAS1068
   BIOS version : 6.06.04.00
   Firmware version : 0.09.03.219
   Channel description : 1 Serial Attached SCSI
   Initiator ID : 63
   Maximum physical devices : 62
   Concurrent commands supported : 511
   Slot : 6
   Bus : 3
   Device : 1
   Function : 0
   RAID Support : Yes
------------------------------------------------------------------------
IR Volume information
------------------------------------------------------------------------
IR volume 1
   Volume ID : 7
   Status of volume : Okay (OKY)
```

```
         RAID level : 1
         Size (in MB) : 34332
         Physical hard disks (Target ID) : 11 8
      ------------------------------------------------------------------
      Physical device information
      ------------------------------------------------------------------
      Initiator at ID #63
      Target on ID #8
         Device is a Hard disk
         Slot # : 2
         Target ID : 8
         State : Online (ONL)
   Size (in MB)/(in sectors) : 34732/71132960
   Manufacturer : HP
   Model Number : DG036A8B5B
   Firmware Revision : HPD1
   Serial No : B2G1P51003CD0501
Target on ID #9
   Device is a Hard disk
   Slot # : 3
   Target ID : 9
   State : Ready (RDY)
   Size (in MB)/(in sectors) : 34732/71132960
   Manufacturer : HP
   Model Number : DG036A8B53
   Firmware Revision : HPD3
   Serial No : 3LC00NLK000085159S8A
Target on ID #10
   Device is a Enclosure services device
   Slot # : 10
   Target ID : 10
   State : Standby (SBY)
   Manufacturer : HP
   Model Number : MSA50 -10D25G1
   Firmware Revision : 1.20
   Serial No :
Target on ID #11
   Device is a Hard disk
   Slot # : 1
   Target ID : 11
   State : Online (ONL)
   Size (in MB)/(in sectors) : 34732/71132960
   Manufacturer : HP
   Model Number : DG036A8B53
   Firmware Revision : HP53
   Serial No : 3LC01PSA00008515A1Y7
------------------------------------------------------------------
Enclosure information
------------------------------------------------------------------
```

```
Enclosure# : 1
   Logical ID : 500605b0:0000f5d0
   Numslots : 7
   StartSlot : 2
   Start TargetID : 0
   Start Bus : 0
Enclosure# : 2
   Logical ID : 500508b3:00a0535c
   Numslots : 11
   StartSlot : 0
   Start TargetID : 0
   Start Bus : 0
     ------------------------------------------------------------
```

TABLE 6-1 Logical drive status values:

Okay (OKY)	Volume is Active and drives are functioning properly. User data is protected if the volume is IM or IME
Degraded (DGD)	Volume is Active. User data is not fully protected due to a configuration change or drive failure
Rebuilding (RBLD)	Data resync or rebuild may be in progress
Inactive, Okay (OKY	Volume is inactive and drives are functioning properly. User data is protected if the current RAID level is RAID 1 (IM) or RAID 1E (IME).
Inactive, Degraded (DGD)	Volume is inactive and the user's data is not fully protected due to a configuration change or drive failure; a data resync or rebuild may be in progress.

TABLE 6-2 Physical device status values:

Online (ONL)	Drive is operational and is part of a logical drive.
Hot Spare (HSP)	Drive is a hot spare that is available for replacing a failed drive in an array
Ready (RDY)	Drive is ready for use as a normal disk drive; or it is available to be assigned to a disk array or hot spare pool

Available (AVL)	Drive may or may not be ready, and it is not suitable for inclusion in an array or hot spare pool (i.e., it is not spun up, its block size is incorrect, or its media is removable)
Failed (FLD)	Drive was part of a logical drive or was a hot spare drive, and it failed. It has been taken offline
Standby (SBY)	This status is used to tag all non-hard drive devices.

`hotspare` command

The `hotspare` command creates a hot spare disk drive, which is added to hot spare pool 0. The number of disk drives in an IM, IME, or IS volume, including the hot spare disk cannot exceed six for LSI53C1020/1030 controllers and eight for LSI1064/1064E and LSI1068/1068E controllers. Only one hot spare disk can be created. The capacity of the hot spare disk must be greater than or equal to the capacity of the smallest disk in the logical drive. An easy way to verify this is to use the `display` Command.

The following rules must be observed when creating hot spare disks:

- A hot spare disk cannot be created unless at least one IM or IME volume is already created
- For LSI1064/1064E and LSI1068/1068E controllers, CFGGEN does not allow adding a hot spare disk of a type (SAS/SATA) that is different from the disk types in any of the volume

Example

```
cfggen controller-number hotspare <SCSI ID> [delete]

cfggen controller-number hotspare <enclosure:bay> [delete]
```

controller-number	Number of the SCSI bus or SAS controller targeted by this command
SCSI ID	SCSI ID of the drive targeted by this command
enclosure:bay	The enclosure:bay value for the disk drive to use for the new hotspare disk. These values can be obtained via the output of the DISPLAY command
[delete]	Optional delete action to specify that the hot spare drive with<SCSI ID> needs to be deleted.

Program Return Value

0x00 SUCCESS command completed successfully
0x01 FAILURE bad command line arguments or operational failure
0x02 ADAPTER_NOT_FOUND adapter specified cannot be found

`list` command

The LIST command displays a list of all controllers present in the system, along with their corresponding controller #.

Example

```
cfggen list
```

Parameters

None

Program Return Value

```
0x00 SUCCESS command completed successfully
0x01 FAILURE bad command line arguments or operational failure
0x02 ADAPTER_NOT_FOUND adapter specified cannot be found
```

Sample Output

Here is an example of the output of LIST command

Note – The format and fields in the output may vary on different versions):

```
Adapter Vendor Device Segment SubSys SubSys
Index Type ID ID ID Bus Device Func Ven ID Dev ID
----- --------- ------ ------ -------- --- ------ ---- ------ ------
0 53C1030 1000h 30h 0000h 02h 03h 00h 0e11h 00dah
1 53C1030 1000h 30h 0000h 02h 03h 01h 0e11h 00dah
2 SAS1068 1000h 54h 0000h 09h 02h 00h 1000h 3050h
```

`rebuild` command

The REBUILD command initiates a resync of drives in an IM or IME volume. This command is used to force a manual resync of drives in the volume even if the auto rebuild is turned off. This command is accomplished by bringing the secondary drive offline and bringing it online immediately there by kicking a resync. The volume status changes to Resyncing (RSY) upon successful execution.

Example

cfggen <controller #> **rebuild** <volume id>

controller-number	Number of the controller targeted by this command
volume id	A valid SCSI Id of a volume of type IM or IME

Program Return Value

0x00 SUCCESS command completed successfully
0x01 FAILURE bad command line arguments or operational failure
0x02 ADAPTER_NOT_FOUND adapter specified cannot be found

`status` Command

The `status` command displays the status of any volume synchronization operation that is currently in progress on the controller. If no such operation is in progress, cfggen displays a message indicating this before it exits. The `status` command adds the flag Inactive to the Volume State field, if the controller firmware marks the volume as Inactive.

Example

`cfggen` *controller-number* **`status`**

controller-number	Number of the SCSI bus or SAS controller targeted by this command

Program Return Value

```
0x00 SUCCESS command completed successfully
0x01 FAILURE bad command line arguments or operational failure
0x02 ADAPTER_NOT_FOUND adapter specified cannot be found
```

Sample Output

Here is an example of the status information returned when a volume resynchronization is in progress:

Background command progress status for controller 0...
IR Volume 1
Current operation : Synchronize
Volume ID : 6
Volume status : Enabled
Volume state : Degraded
Physical disk I/Os : Not quiesced
Volume size (in sectors) : 70311936
Number of remaining sectors : 68250624
Percentage complete : 2.93%

Here is an example of the status information returned when no background volume operation is in progress:

Background command progress status for controller 0...
IR Volume 1
Current operation : None
Volume ID : 6
Volume status : Enabled
Volume state : Optimal
Physical disk I/Os : Not quiesced

The status fields in the data displayed can have the following values:

Current operation - Synchronize or None
Volume status - Enabled or Disabled
Volume state - [Inactive] Optimal, Degraded or Failed
Physical disk I/Os - Quiesced or Not quiesced

Monitoring and Managing RAID Arrays

This section describes how to perform common task, such as create, rebuild, and monitor RAID0 and RAID1 arrays.

You must determine the controller numbers used. The controllers are enumerated starting with 0 based on bus location. Unless other LSI add-on cards have been installed, the controller number for the 1064 is 0. Otherwise, run the `MPTutil` to determine the order of the LSI controllers.

This section describes the following tasks:

- "To Create a RAID 0 Array" on page 53
- "To Fail RAID 0" on page 55
- "To Create a RAID 1 Array" on page 56
- "To Rebuild a RAID 1 Array" on page 56
- "To Delete a RAID Array" on page 57

▼ To Create a RAID 0 Array

1. **Check that no other arrays have been created for the desired controller with the `status` command.**

```
> cfggen controller-number status
```

2. **If there is an array present, delete the array as described in** "To Delete a RAID Array" on page 57**.**

3. **Determine the slot numbers of the desired drives and check the drives are ready with the display command.**

```
> cfggen controller-number display
```

This command gives information about the controller, IR volume, physical devices, and enclosure. The slot numbers are located in the physical device information for each device. For RAID 0, at least two slot numbers are needed. An example of the physical device information is below.

The slot number and drive state are in **bold**:

```
Target on ID #1
Device is a Hard disk
Slot # : 1
Target ID : 11
State : Ready (RDY)
```

4. **Create RAID0 array by using the** `create` **command:**

```
> cfggen controller-number create IS size slot-numbers [noprompt]
```

The size is in Mbytes and, like the `delete` command, `noprompt` is optional. For example, to create a RAID0 array on controller 0 that is 512 MB on slots 0 and 1, type:

```
> cfggen 0 create IS 512 0 1
```

If an array already exists on the specified drives, the `create` command gives an error stating there are not enough resources.

5. **Rerun the** `status` **command and verify that the new RAID0 volume is detected.**

The command displays:

```
Current Operation: None
Volume Status: Enabled
Volume State: Optimal
Physical Disk I/Os: Not Quiesced
```

▼ To Fail RAID 0

Make RAID0 fail by removing one of its drives.

1. **Create a RAID0 array.**
2. **Remove one of the drives.**
3. **Check that RAID0 fails by executing the `status` command.**
   The command displays:

```
Current Operation: None
Volume Status: Enabled
Volume State: Failed
Physical Disk I/Os: Not Quiesced
```

4. **Replace the drive with a new drive.**
5. **Verify that the drive still fails by executing the `status` command.**
   The command displays:

```
Current Operation: None
Volume Status: Enabled
Volume State: Failed
Physical Disk I/Os: Not Quiesced
```

6. **Delete the array as described in** "To Delete a RAID Array" on page 57**.**
7. **Create RAID0 as described in** "To Create a RAID 0 Array" on page 53**.**
8. **Once the drive is created, check the status.**
   The command displays:

Current Operation: None
Volume Status: Enabled
Volume State: Optimal
Physical Disk I/Os: Not Quiesced

▼ To Create a RAID 1 Array

Delete current RAID arrays and determine slot numbers for the desired drives as described in "To Delete a RAID Array" on page 57. For a RAID1 array, you need to specify two slot numbers.

1. **Create the RAID1 array.**

```
> cfggen controller-number create IM size slot-numbers [noprompt]
```

2. **Verify the RAID1 has been created by running the `status` command.**

   The command displays:

```
Current Operation: None
Volume Status: Enabled
Volume State: Optimal
Physical Disk I/Os: Not Quiesced
```

▼ To Rebuild a RAID 1 Array

1. **Remove a drive.**
2. **Run the `status` command.**

   The command displays:

```
Current Operation: None
Volume Status: Enabled
Volume State: Degraded
Physical Disk I/Os: Not Quiesced
```

3. **Insert the drive.**
4. **Run the `display` command.**

   The drive should have a volume state of Out of Sync.
5. **Run the `status` command.**

   Depending on the size of the volumes, the status should be in one of two states:

   - Synchronized

- Done rebuilding

If the drive is still in process of rebuilding, the command displays:

```
Current Operation: Synchronized
Volume Status: Enabled
Volume State: Degraded
Physical Disk I/Os: Not Quiesced
Percentage complete: x.xx%
```

If the drive is finished being rebuilt, the command displays:

```
Current Operation: None
Volume Status: Enabled
Volume State: Optimal
Physical Disk I/Os: Not Quiesced
```

▼ To Delete a RAID Array

1. **Verify that there is a RAID array present by running the `status` command.**
2. **Reboot the system.**
3. **Delete the RAID array.**

> cfggen *controller-number* delete [noprompt]

If you use the `noprompt` option, the utility automatically deletes the arrays. Otherwise, the utility asks if it can continue with this command.

4. **Verify that the array is deleted by running the `status` command.**

   The command reports that there are no volumes present.

5. **Reboot the system.**

PART III MegaRAID Storage Manager

This part describes how to use the MegaRAID Storage Manager and has the following chapters:

- "MegaRAID Storage Manager (MSM) Installation" on page 61
- "Using MegaRAID Storage Manager" on page 67
- "LSI SNMP Utility" on page 99

CHAPTER 7

MegaRAID Storage Manager (MSM) Installation

The MegaRAID Storage Manager (MSM) program provides you with graphical user interface (GUI) tools to configure RAID storage systems, based on the LSI 106*x* controllers used in some of the x64 servers. To determine if your server supports this program, refer to the Product Notes for your platform.

This chapter includes the following topics:

- "Overview" on page 61
- "Installing the Program" on page 62

Note – MegaRAID Storage Manager (MSM) is also known as MegaRAID Storage Manager - Integrated RAID (MSM-IR) for versions earlier than 2.*x*. Some x64 servers support MSM 2.*x*; others support MSM 1.*x* versions.

Overview

The MSM program enables you to configure the controllers, physical disk drives, and virtual disk drives on your system. The Configuration Wizard in the MSM program simplifies the process of creating disk groups and virtual disk drives by guiding you through several simple steps to create your storage configurations.

MSM works with the appropriate operating system (OS) libraries and drivers to configure, monitor, and maintain storage configurations attached to x64 servers. The MSM GUI displays device status in the form of icons, which represent the controllers, virtual disk drives, and physical disk drives on your system. Special

icons appear next to the device icons on the screen to notify you of disk failures and other events that require immediate attention. System errors and events are recorded in an event log file and are displayed on the screen.

Installing the Program

You can install MegaRAID Storage Manager-IR on x64 Servers with the Microsoft Windows Server 2003, Red Hat Linux, and SUSE Linux operating systems.

This section includes the following topics:

- "Installing MSM on the Windows OS" on page 62
- "To Install MSM on the Linux OS" on page 63

Installing MSM on the Windows OS

Use this procedure to install the MSM program on an x64 server with Windows Server 2003.

Note – The MSM installation files and drivers were installed on your system if you selected the correct optional components during the Windows 2003 Server installation. If you did not select these components, continue with this procedure.

The MSM packages are available on the product Tools and Drivers CD, and also as part of an archive called `windows.zip`. You can download CD ISO and the archive from the Sun web site. See "Obtaining Utilities" on page xiii.

▼ To Install MSM on the Windows OS

1. **Insert the Tools and Drivers CD into your server's CD-ROM drive. Alternately, you can extract the contents of the `windows.zip` archive.**
2. **Navigate to `\windows\wk3\packages\` (on the CD) or `wk3\packages\` (in the archive).**
3. **Run the installation application in this directory. This is a file with a name of the form `InstallPack`*xxxxxx*`.zip`, where the *xxxxxx* is a version string.**

   The Sun Fire Installation Package dialog box appears.
4. **Click the Optional Components check box.**

5. **Click Next to accept the settings.**
6. **Review the Important Note, and then click Next.**

   The Welcome to the Sun Fire Installation Wizard appears.
7. **Click Next.**

   The End User License Agreement dialog box appears.
8. **Select I accept this agreement, and then click Next.**
9. **Click Finish.**
10. **Click Yes to restart your system to complete the installation.**

    The MSM program has been installed on your system.

Windows Installation Error Messages

The Windows Installer program might display error messages during the installation process. The error message text is self-explanatory. If you need more information about these messages, see the list of installer error messages on the Microsoft Developers Network (MSDN) web site at:

http://msdn.microsoft.com/library/default.asp?url=/library/en-us/msi/setup/windows_installer_error_messages.asp

Installing Windows Drivers

For information about drivers for the Windows Server 2003 Operating System (OS), refer to the *Windows Operating System Installation Guide* for your platform.

▼ To Install MSM on the Linux OS

Use this procedure to install the MSM utility and the required drivers on a Sun x64 Server with the Linux OS.

1. **Make sure the driver `mptctl` is inserted.**

   This module is required for a Linux system to be administered by MSM. To find out if the module is correctly inserted, type the command:

   ```
   # lsmod | grep mptctl
   ```

   If you see output similar to this, then `mptctl` is inserted. If the command has no input, `mptctl` is not inserted.

   To insert `mptctl`, type this command:

   ```
   # modprobe mptctl
   ```

2. **To ensure you have a fully operational system to run the MSM utility, install a full Linux installation (everything) on the system.**

   Using the MSM utility on a Linux system requires a number of shared libraries that are not included in the basic install of most Linux distributions.

3. **Locate the files before proceeding.**

   The files are located on the Tools and Drivers CD provided with your system, if available, or from the Tools and Drivers CD image, downloadable from your product web site. For instructions on downloading the application from the Tools and Drivers CD image, see "Obtaining Utilities" on page xiii.

   The MSM files are in `/linux/tools/raid` directory.

4. **Insert the Tools and Drivers CD into the CD-ROM drive connected to your server, or copy the files to your system.**

5. **Locate the file MSM Linux installer file in the `raid` directory.**

6. **Copy the installer to your home directory by typing:**

   `# cp` *MSM-Linux-installer* **`~`**

7. **Uncompress and then untar the installer**

   `# tar zxvf` *MSM-Linux-installer*

8. **Locate the `disk` directory created by uncompressing the installer, and move to this directory by typing:**

   ```
   # cd disk
   ```

9. **Run the install script by typing:**

   ```
   # ./install.sh
   ```

10. **Read and accept the License agreement.**

11. **Select the type of install that best suits your needs.**

▼ To Create a Partition and Filesystem on Linux OS

1. **Determine which device is the RAID device.**
2. **Check the last entries in the `dmesg` log to determine which SCSI device is the LSI RAID device. For example, type:**

   `# dmesg | tail -30 | grep Attached`

   This searches the previous 30 lines of the `dmesg` for the appropriate line. A line such as the following will appear:

   ```
   Attached scsi disk sda at scsi0, channel 0, id 2 lun 0
   ```

   In this case the disk *sda* would be the SCSI device.
3. **Create a new partition by typing:**

   `# fdisk /dev/`*device*

   a. **Type n to create a new partition.**

   b. **Type p to make it a primary partition.**

   c. **Press enter to choose the default and start the partition at the beginning of the drive.**

   d. **Press enter to choose the default and end the partition at the end of the drive.**

   e. **Type w to write the new partition to disk.**

Note – You can create a partition that is smaller than the maximum size of the device or create multiple partitions on the device. See the `fdisk` man page for details.

---

4. **Create a new file system.**

Note – You can create a number of different types of file systems under Linux. This section covers a default configuration of ext2 and ext3 file systems. Refer to the documentation included with the file system of your choice for complete instructions.

You determined the base name of the RAID device in Step 2 above. That device name refers to the RAID as a whole. When you created a partition on that RAID, a device name was created to address just that partition. If the RAID is sda, and you followed the above directions, the partition device would be sda1. If there were multiple partitions on a disk there would be multiple corresponding device names (sda1, sda2, sda3, etc).

Now that a partition exists on the raid device, a file system needs to be written to that partition. The following commands create an ext2 or ext3 file system. Replace *device* with the device name referencing the appropriate partition.

a. For an ext2 file system, type:

```
# mke2fs -i 1024 /dev/device
```

b. For an ext3 file system, type:

```
# mke2fs -i 1024 -j /dev/device
```

CHAPTER 8

Using MegaRAID Storage Manager

This chapter explains how to launch and use the MSM (MegaRAID Storage Manager) program. Use the program to create RAID arrays, and then manage and monitor the RAID arrays after array creation. The following sections describe how to start and use the MSM program:

- "Starting the MSM Program" on page 67
- "Using the MSM RAID Configuration Wizard" on page 78
- "Monitoring System Events and Storage Devices" on page 90
- "Maintaining and Managing Storage Configurations" on page 94
- "Known Issues" on page 96

Note – MegaRAID Storage Manager (MSM) is also known as MegaRAID Storage Manager - Integrated RAID (MSM-IR) for versions earlier than 2.*x*. Some x64 servers ship MSM 2.*x*; others support MSM 1.*x* versions.

Starting the MSM Program

The following sections describe how to start and log in to the MSM program:

- "To Start MSM on the Windows 2003 Server" on page 68
- "To Start MSM on a RHEL 4 Server" on page 68
- "To Start MSM on a SLES 10 Server" on page 68
- "Running MSM" on page 68
- "To Access Servers on Alternate Subnets" on page 70
- "To Log in to MSM" on page 70

▼ To Start MSM on the Windows 2003 Server

1. **On the taskbar, click Start and choose All Programs.**
2. **Choose the MSM program group, and click StartupUI.**

Tip – Alternatively, double click the MSM icon on the desktop.

▼ To Start MSM on a RHEL 4 Server

- **On the taskbar, click Applications > System Tools > MegaRAID Storage Manager StartupUI.**

▼ To Start MSM on a SLES 10 Server

- **On the taskbar, click Computer > System > MegaRAID Storage Manager StartupUI.**

Running MSM

After you have started MSM, the MSM server window appears. The first screen is the Select Server window, similar to the one shown in FIGURE 8-1.

Note – If a warning appears indicating the Windows Firewall has blocked some features of the program, click **Unblock** to allow MSM to start. The Windows Firewall might block some Java based programs like MSM. If there are multiple servers on the network, you might experience a delay before the Select Server window appears. A network with multiple servers might look similar to the Select Server window shown in FIGURE 8-1.

FIGURE 8-1 MSM Select Server Window (Opening Screen)

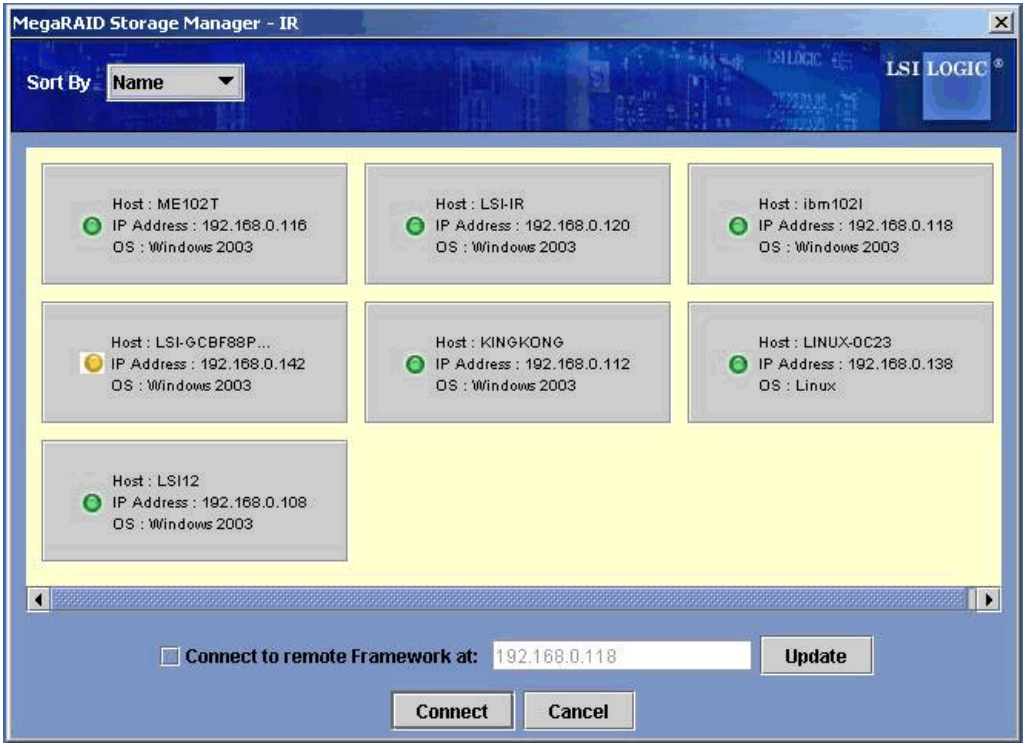

The Host server icon status is indicated by an LED-like indicator located within the host server icon, to the left of the center to the left of the IP address.

- A green LED indicates normal operation.
- A yellow LED indicates that the server is running in a degraded state.

  For example, a disk drive used as a virtual disk has failed.
- A red LED indicates that the server's storage configuration has failed.

Note – You can access servers on a different subnet by entering an IP address in the *Connect to remote Framework at* field at the bottom of the screen. The check box next to the *Connect to remote Framework at* field enables you to access a standalone server running MSM, if it has a network connection.

▼ To Access Servers on Alternate Subnets

1. **Type the IP address in the** `Connect to remote Framework at` **field, located at the bottom of the screen.**

   Use the IP address of a server operating in the desired subnet where MSM has access.

2. **Click Update.**

3. **To access the standalone server running MSM with a network connection, select the check box to the left of** `Connect to remote Framework at`**.**

▼ To Log in to MSM

1. **Double-click the icon of the desired Host server in the** `Select Server` **window.**
   See FIGURE 8-1 **or** FIGURE 8-2**.**

FIGURE 8-2 MSM Select Server Window (Simplified)

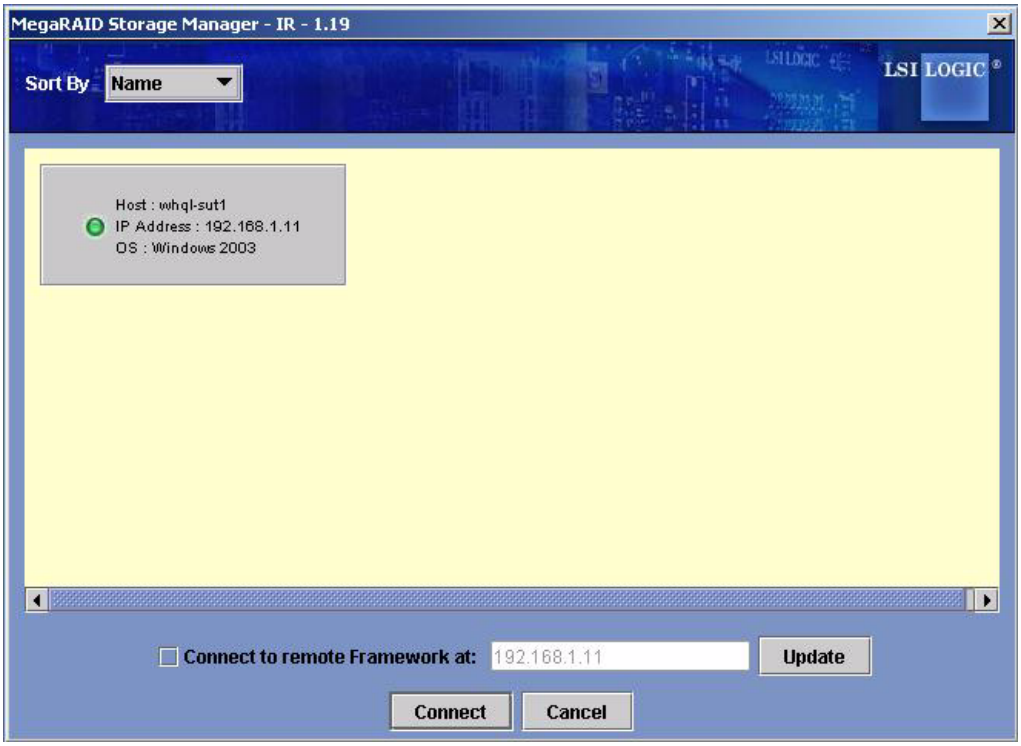

2. **Select the check box next to** `Connect to remote Framework at` **in the lower portion of the screen, as shown in** FIGURE 8-2**.**

   The User Name & Password dialog box appears. See FIGURE 8-3.

FIGURE 8-3 MSM User Name & Password Dialog Box

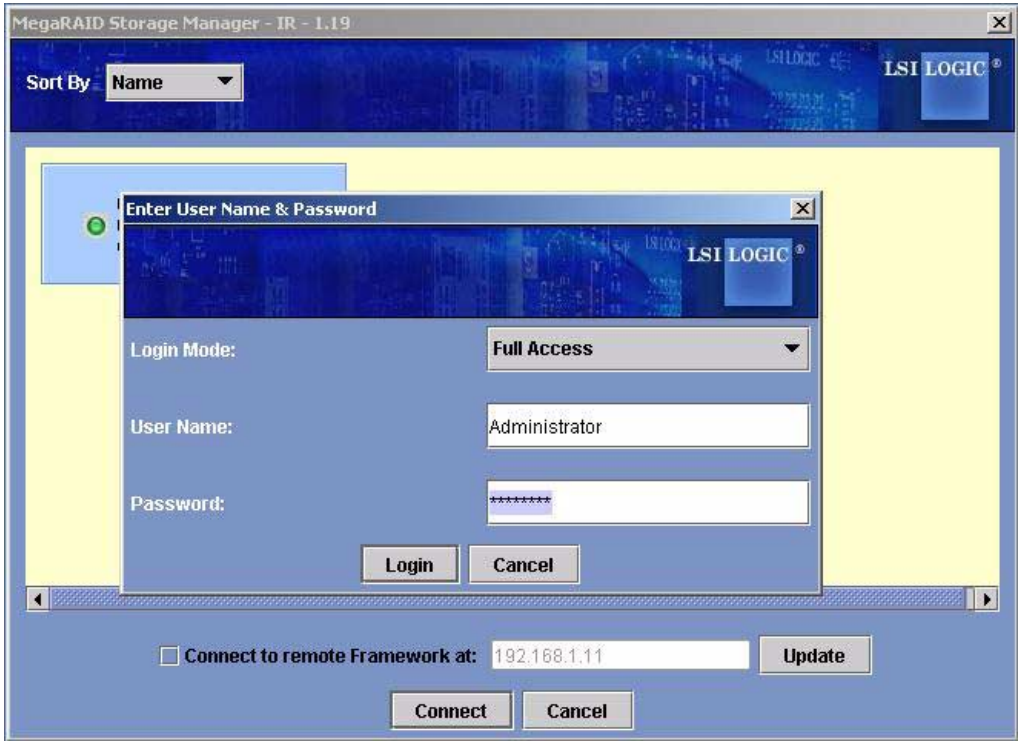

3. **Select a login mode from the drop-down list. See** FIGURE 8-3**.**
   - Select **Full Access** to view or modify the current configuration.
   - Select **View Only** to view and monitor the configuration.

Note – If you are accessing the server over a network, you will also need to enter the root/administrator user name and password to use Full Access mode. Step 4 gives you access to the server, but not full access over the network.

If your user name and password are correct, the MSM Physical/Logical window appears, similar to the one shown in FIGURE 8-4.

4. **(View Only Access) Type your user name and password, and then click Login.**

5. **(Full Access) Type the Administrator user name and password, and then click Login. See** FIGURE 8-3**.**

   The MSM Physical/Logical window displays similar to the one shown in FIGURE 8-4.

FIGURE 8-4 MSM Physical/Logical Window

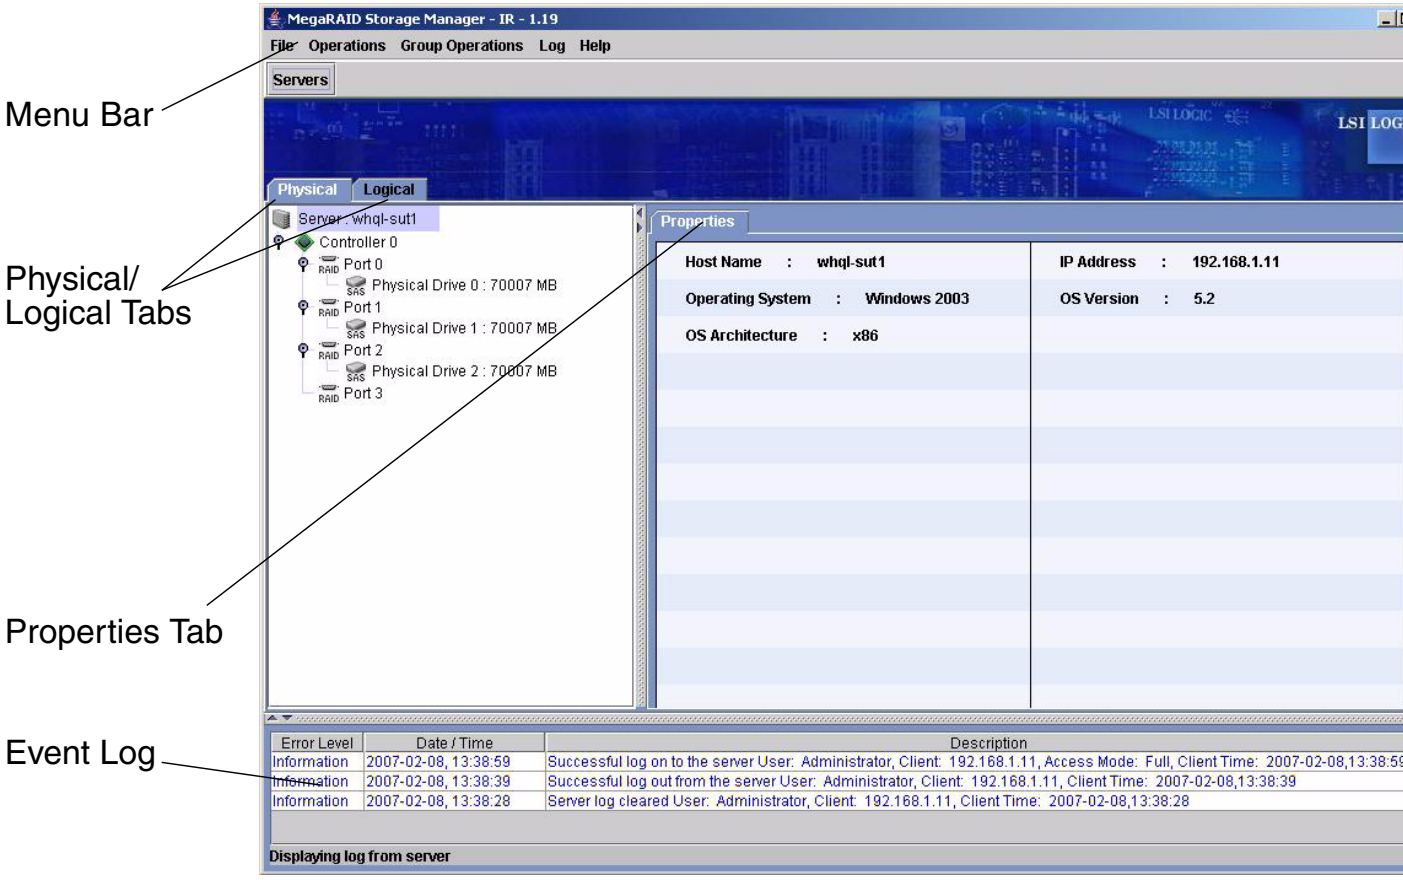
Menu Bar
Physical/
Logical Tabs
Properties Tab
Event Log
MegaRAID Storage Manager - IR - 1.19
File Operations Group Operations Log Help
Servers
Physical
Logical
Server : whql-sut1
Controller 0
Port 0
Physical Drive 0 : 70007 MB
Port 1
Physical Drive 1 : 70007 MB
Port 2
Physical Drive 2 : 70007 MB
Port 3
Properties
Host Name : whql-sut1
IP Address : 192.168.1.11
Operating System : Windows 2003
OS Version : 5.2
OS Architecture : x86
Error Level
Date / Time
Description
Information 2007-02-08, 13:38:59 Successful log on to the server User: Administrator, Client: 192.168.1.11, Access Mode: Full, Client Time: 2007-02-08,13:38:59
Information 2007-02-08, 13:38:39 Successful log out from the server User: Administrator, Client: 192.168.1.11, Client Time: 2007-02-08,13:38:39
Information 2007-02-08, 13:38:28 Server log cleared User: Administrator, Client: 192.168.1.11, Client Time: 2007-02-08,13:38:28
Displaying log from server

MSM Windows

This section describes the MSM Physical/Logical window, which appears after you have logged in to the MSM program. The following topics describe the panels and menu options that display in this window.

Physical/Logical View Panel

The left panel of the MSM window displays either the *Physical view* or the *Logical view* of the host server and the devices, depending on which tab you select.

- Physical view: Shows the hierarchy of physical devices in the host server. At the top of the hierarchy is the host server. The host server has one or more controllers installed and each controller has one or more ports. Disk drives and other devices are attached to the ports.
- Logical view: Shows the hierarchy of controllers, virtual disks, and disk groups that are defined on the host server. (Physical drives also appear in the Logical view, so you can see which physical disk drives are used by each virtual disk.)

TABLE 8-1 shows the icons that appear in the left panel to represent the controllers, disk drives, and other devices:

TABLE 8-1 MSM Window Icons

System	Degraded state -- A yellow LED indicates that the device is running in a degraded state. For example, the yellow LED next to virtual disk indicates that it is running in a degraded state because a disk drive has failed.
Controller	
Port	

TABLE 8-1 MSM Window Icons

Disk group	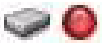 Device failure -- A red LED indicates that the device has failed. For example, the red LED next to disk drive icon indicates that it has failed.
Virtual disk	
Physical disk drive	

Event Log Panel

The lower part of the MSM window displays the system event log entries. New event log entries appear during the session. Each entry has a timestamp and date, an error level indicating the severity of the event, and a brief description of the event.

Note – Some events in MSM_IR do not display an accurate date and timestamp value in the MSM log. When this is the case, the date and timestamp line will display "#####" as the value.

For example, creating a new RAID 0 or RAID 1 will generate "#####" in the MSM date & timestamp log. Also, swapping hard disks will only display "####################" in the MSM log section.

FIGURE 8-5 MSM Window (Events and Icons)

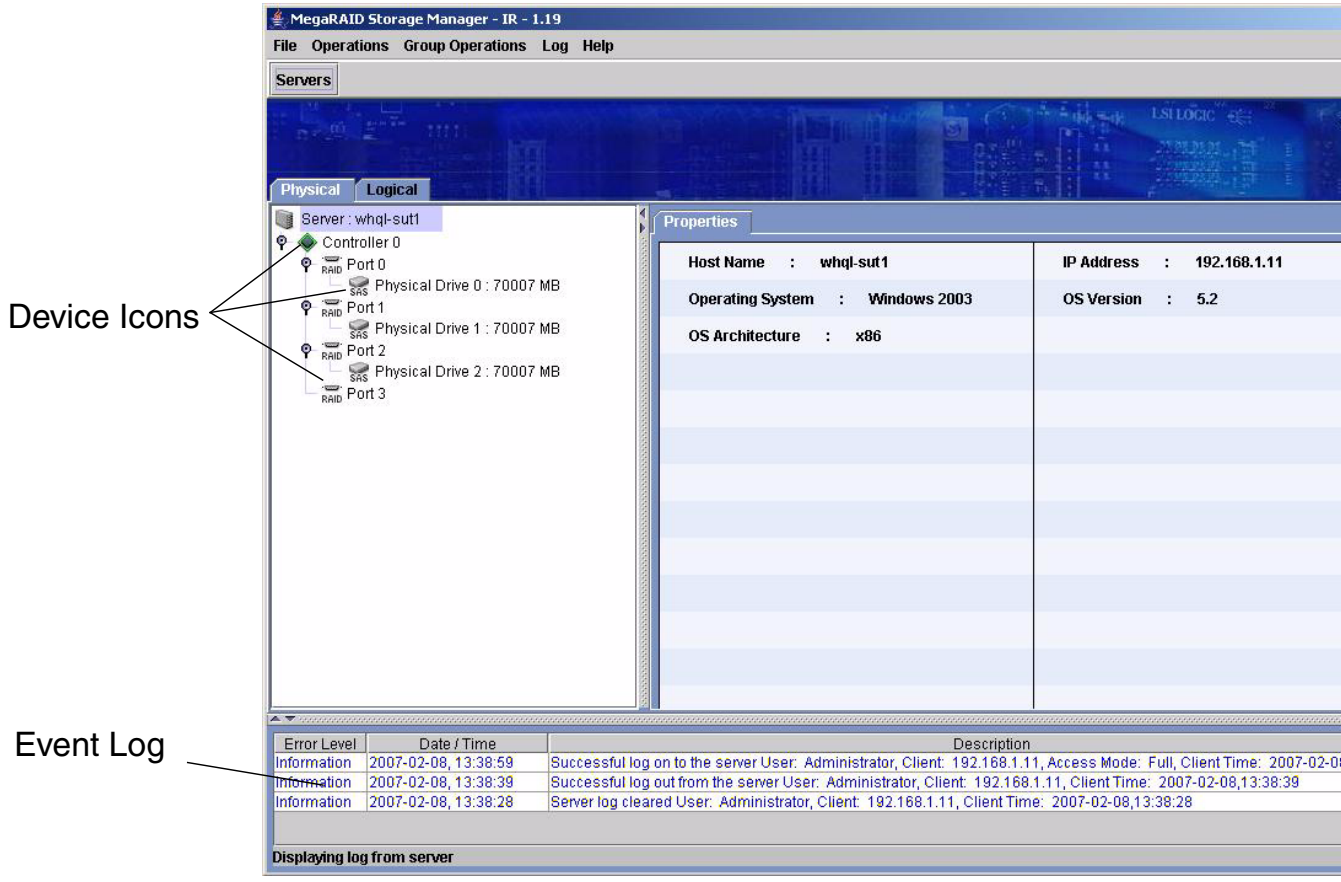
Device Icons
Event Log
MegaRAID Storage Manager - IR - 1.19
File Operations Group Operations Log Help
Servers
Physical Logical
Server : whql-sut1
Controller 0
Port 0
Physical Drive 0 : 70007 MB
Port 1
Physical Drive 1 : 70007 MB
Port 2
Physical Drive 2 : 70007 MB
Port 3
Properties
Host Name : whql-sut1
IP Address : 192.168.1.11
Operating System : Windows 2003
OS Version : 5.2
OS Architecture : x86
Error Level
Date / Time
Description
Information 2007-02-08, 13:38:59 Successful log on to the server User: Administrator, Client: 192.168.1.11, Access Mode: Full, Client Time: 2007-02-08
Information 2007-02-08, 13:38:39 Successful log out from the server User: Administrator, Client: 192.168.1.11, Client Time: 2007-02-08,13:38:39
Information 2007-02-08, 13:38:28 Server log cleared User: Administrator, Client: 192.168.1.11, Client Time: 2007-02-08,13:38:28
Displaying log from server

Properties/Operations/Graphical View Panel

The right panel of the MSM window has two or three tabs, depending on the device type selected in the left panel.

- Properties tab: Displays information about the selected device.
- Operations tab: Lists the operations that can be performed on the device selected in the left panel.

- Graphical tab: Selected in the right panel if a physical drive or a virtual disk is selected in the left panel.
  - In Graphical View, the device's storage capacity is color coded according to the legend shown on the screen.

Dual-Path SAS Disk Drives

MSM version 2.63 and later handles dual-path disk drives automatically.

When you have SAS disk drives with dual paths, you see a single disk in the left panel of the main menu screen with the Physical tab chosen.

Select a drive in this panel and choose the Properties tab in the right panel of the screen. If the disk drive has two paths, you see a SAS Address 0 and SAS Address 1 in the Properties tab. You also see that the Redundant Paths property has the value 'Yes.'

If you remove one of the paths (for example by removing a Multi-Function NEM that connects a server blade to a disk blade), you see only one SAS address and the Redundant Paths property has the value 'No.' When you restore the path, the Redundant Paths property has the value 'Yes' once again.

Note – You can view the Redundant Paths property when you remove and restore a path to verify that your version of MSM is multi-path aware.

You get an Alert Event Notification whenever a second path is added or deleted. The messages are:

- Redundant path inserted.
- Redundant path broken.

Menu Bar

The brief descriptions listed here refer to the main selections in the MSM menu bar.

- File Menu: The File menu has an Exit option for exiting MSM. It also has a Rescan option for updating the display in the MSM-IR window. (Rescan is seldom required; the display normally updates automatically.)
- Operations Menu: The Operations menu is available when a controller, physical drive, or logical drive is selected in the MSM window. The Operations menu options vary, depending on what type of device is selected in the left panel of the MSM window.

- You can also view the Operations tab in the right panel. If an operation requires user inputs before it can be executed, it appears in the Operations tab but not in the Operations menu. A device-specific Operations menu pops up if you right-click a device icon in the left panel.

  An Advanced Operations submenu is also available. This is where you access the Configuration Wizard and other configuration-related commands. To access this menu, select **Operations > Advanced Operations**.
- Group Operations Menu: The Group Operations menu selections include Initialize and Show Progress.
- Log Menu: The Log menu includes selections for saving and clearing the message log.
- Help Menu: On the Help menu you can select **Help > Help** to view the MSM online Help file. You can select **Help > About** to view version information for the MSM.

Note – When you use MSM online Help, you might see a warning message that Internet Explorer has restricted the file from showing active content. If this warning displays, click the active content warning bar and enable the active content.

Using the MSM RAID Configuration Wizard

Use the MSM RAID configuration wizard to create, delete, or modify RAID configurations**.**

Note – Physical disk drives with bootable partitions cannot be used to create a virtual drive on a SAS IR controller with MSM. Drives with bootable partitions do not appear on the list of available drives to create a new virtual drive. To make physical disk drives with bootable partitions available for use in creating a virtual drive, you must clear the bootable flag in the partition or remove the partition.

▼ To Start the MSM RAID Configuration Wizard

1. **In the Physical tab, click Controller 0. See** FIGURE 8-6**.**

   You can also use the Menu bar to access the RAID configuration wizard, by choosing **Operations > Advanced Operations**.

FIGURE 8-6 MSM Window (RAID Creation)

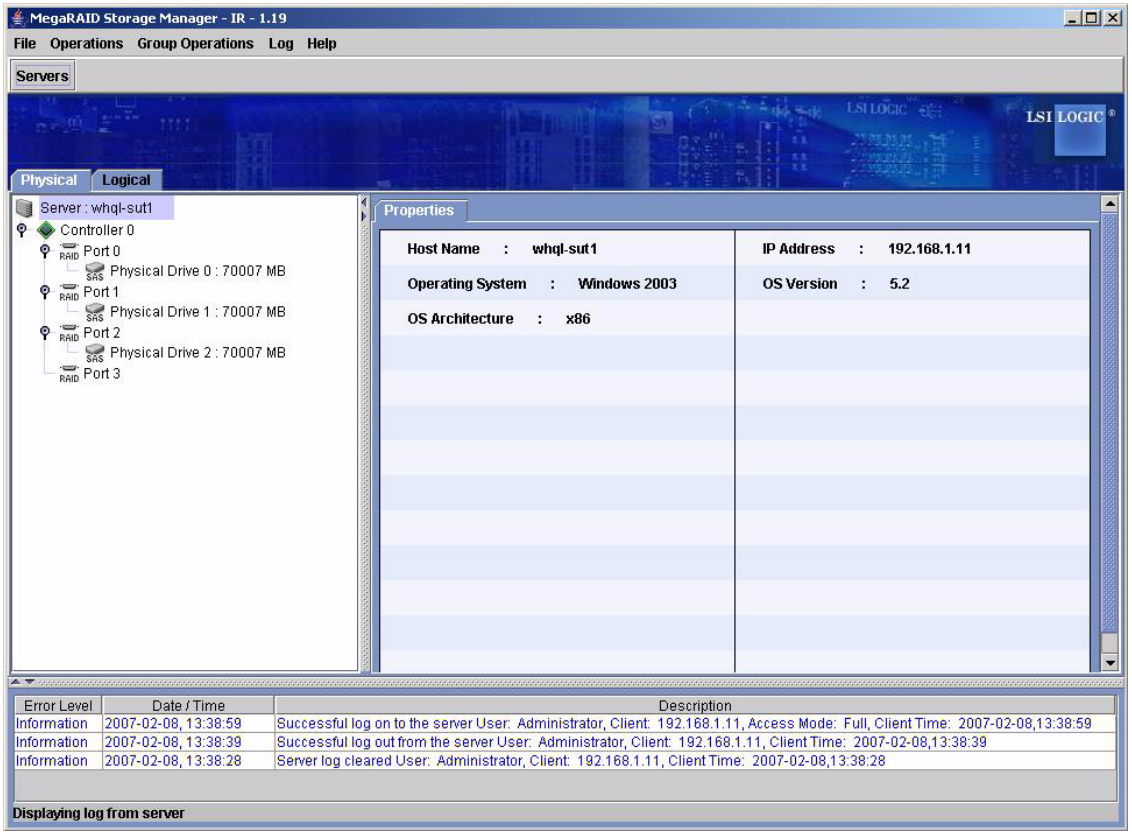
MegaRAID Storage Manager - IR - 1.19
File Operations Group Operations Log Help
Servers
LSI LOGIC
Physical Logical
Server : whql-sut1
Controller 0
Port 0
Physical Drive 0 : 70007 MB
Port 1
Physical Drive 1 : 70007 MB
Port 2
Physical Drive 2 : 70007 MB
Port 3
Properties
Host Name : whql-sut1
IP Address : 192.168.1.11
Operating System : Windows 2003
OS Version : 5.2
OS Architecture : x86
Error Level
Date / Time
Description
Information 2007-02-08, 13:38:59 Successful log on to the server User: Administrator, Client: 192.168.1.11, Access Mode: Full, Client Time: 2007-02-08,13:38:59
Information 2007-02-08, 13:38:39 Successful log out from the server User: Administrator, Client: 192.168.1.11, Client Time: 2007-02-08,13:38:39
Information 2007-02-08, 13:38:28 Server log cleared User: Administrator, Client: 192.168.1.11, Client Time: 2007-02-08,13:38:28
Displaying log from server

FIGURE 8-7 MSM Configuration Wizard Dialog Box

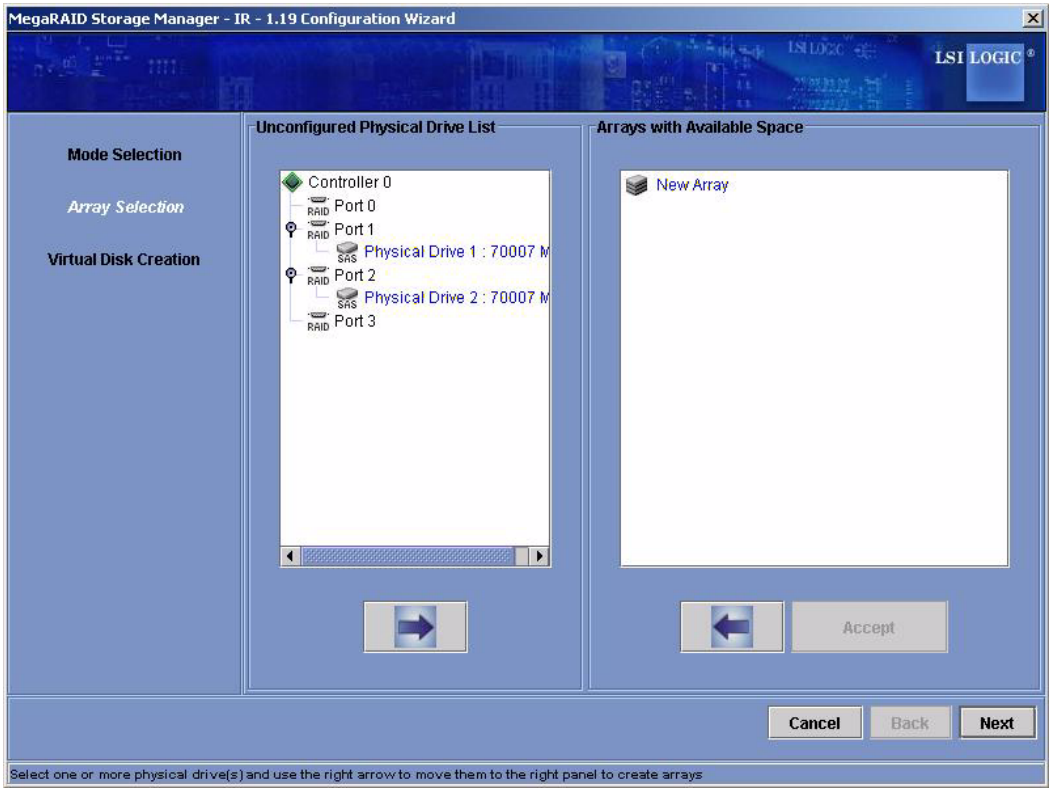

2. **Click Configuration Wizard to start the MSM RAID configuration wizard.**

   The Configuration Wizard dialog box displays.

3. **Select the physical disk drive(s) and add the selected disk drive(s) to the new array.**

Note – Two disks can be selected at the same time by using the shift key after making the first disk drive selection.

---

   a. **In the** `Unconfigured Physical Drive List` **section, click a physical disk drive to select it (highlight).**

   b. **Skip to Step *c*, or make a second selection by holding down the shift key, while your first physical disk drive is selected (highlighted).**

      Both physical disk drives should be highlighted.

c. **Click the right arrow to move the physical disk drive(s) into the section labeled:** `Arrays with Available Space`**.**

4. **When you are finished, click Accept, and then click Next.**

   The Virtual Disk Creation dialog box displays. See FIGURE 8-8.

FIGURE 8-8 MSM Virtual Disk Creation Dialog Box

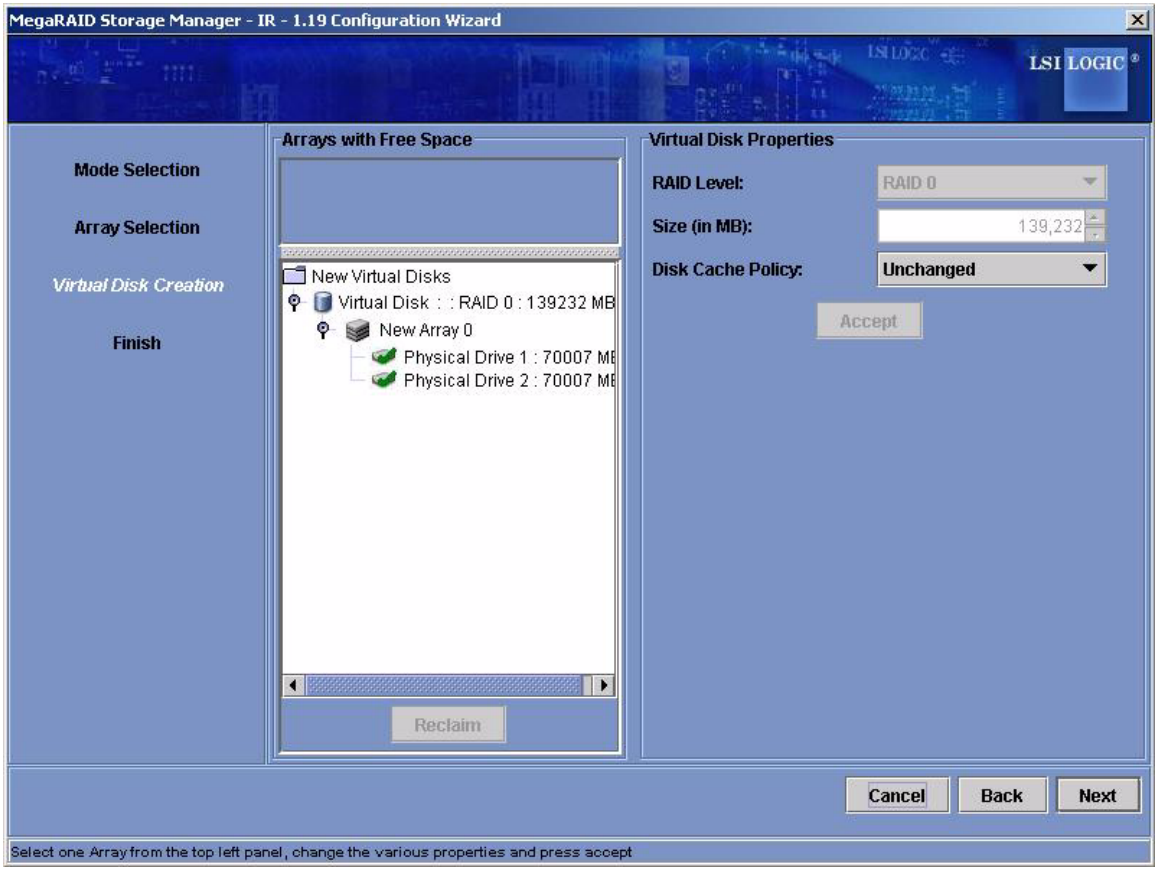

5. **Select a RAID Level, Size (in Mbytes), and a Disk Cache Policy from** `Virtual Disk Properties`**.**

6. **Click the active Accept button, and then click Next. See** FIGURE 8-8 **and** FIGURE 8-9**.**

   The Finish dialog box displays with RAID 0 (New Array) selected and accepted.

FIGURE 8-9 MSM RAID Finish Screen

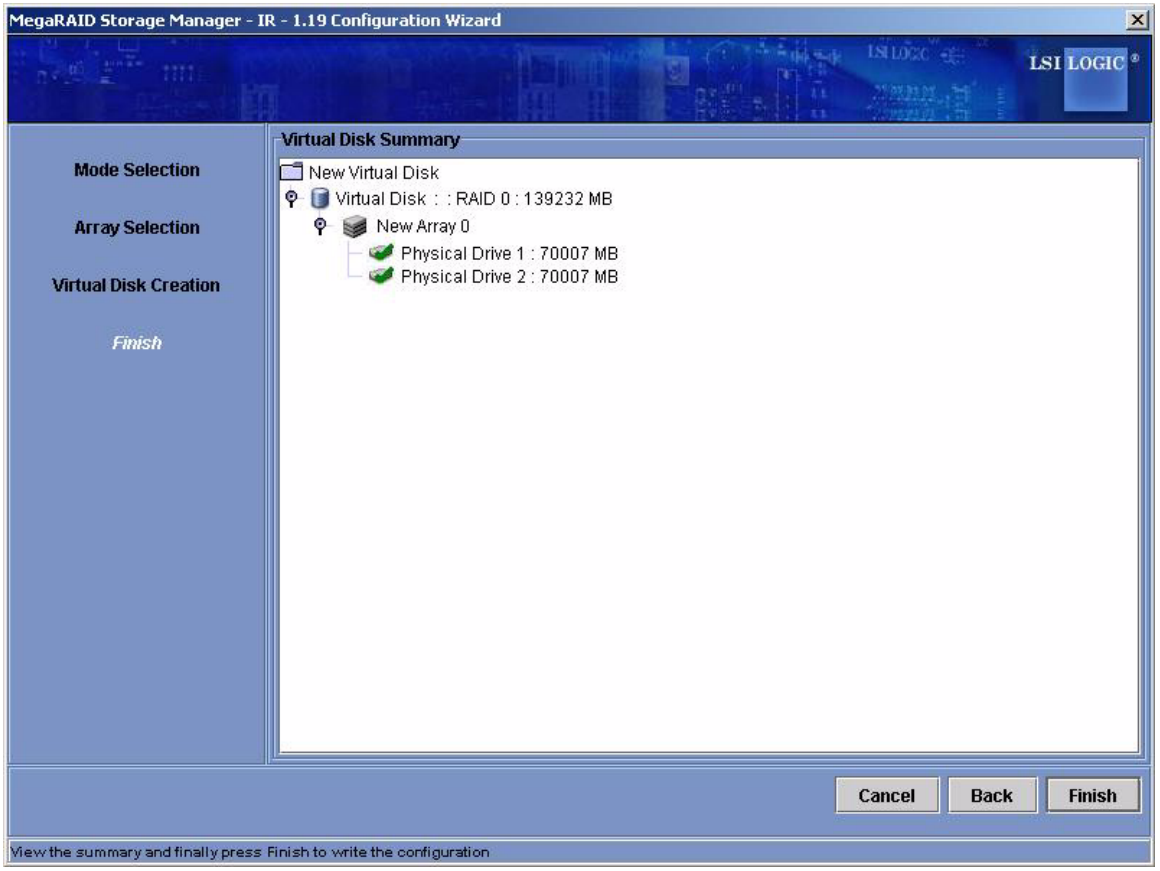

7. Click Finish to create the RAID.

The MSM Configuration Wizard builds the RAID and the resultant updated RAID configuration is displayed in the MSM window. See FIGURE 8-10.

FIGURE 8-10 MSM Physical Window (New RAID Array)

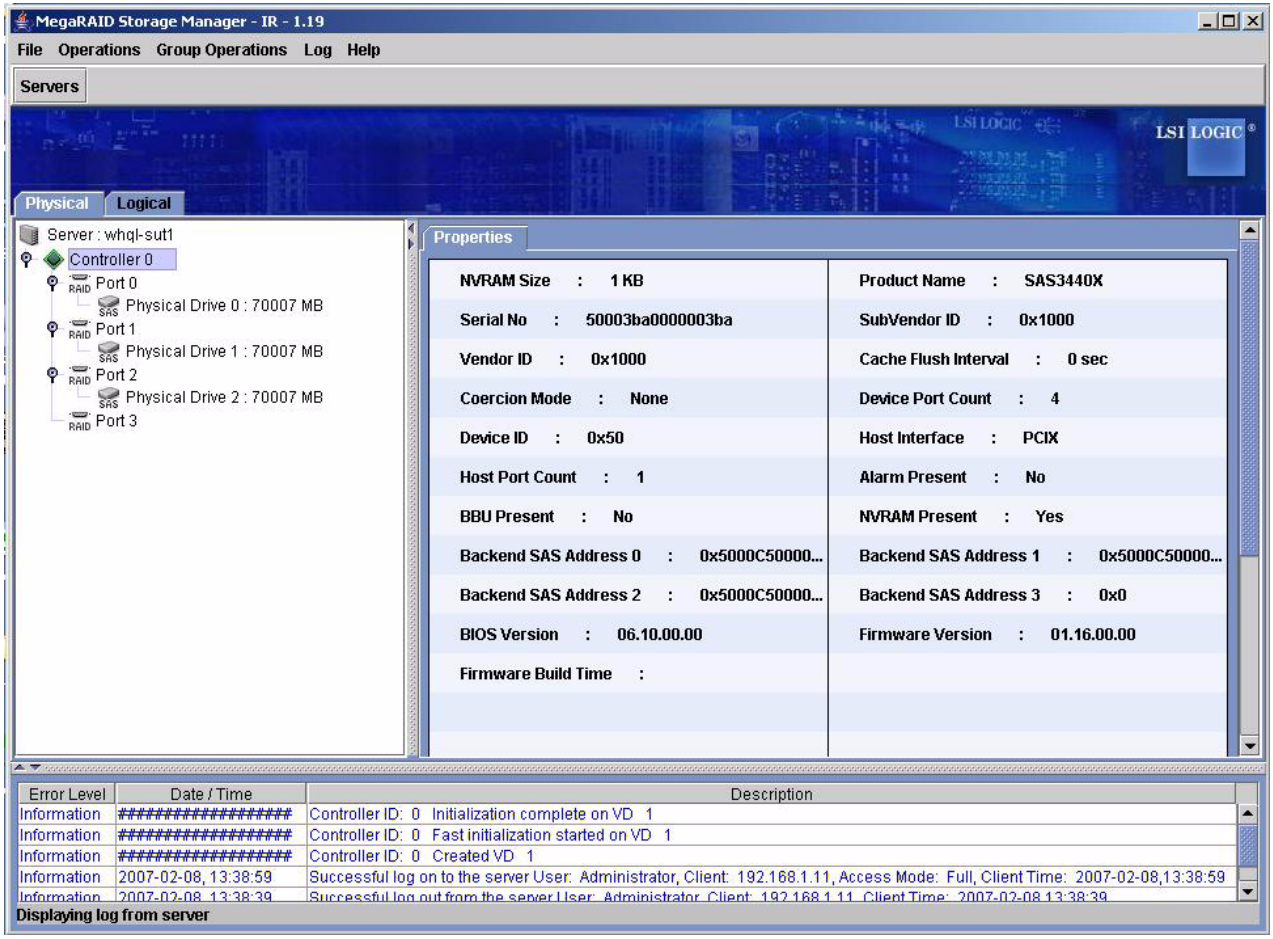
MegaRAID Storage Manager - IR - 1.19
File Operations Group Operations Log Help
Servers
LSI LOGIC
Physical Logical
Server : whql-sut1
Controller 0
Port 0
Physical Drive 0 : 70007 MB
Port 1
Physical Drive 1 : 70007 MB
Port 2
Physical Drive 2 : 70007 MB
Port 3
Properties
NVRAM Size : 1 KB
Serial No : 50003ba0000003ba
Vendor ID : 0x1000
Coercion Mode : None
Device ID : 0x50
Host Port Count : 1
BBU Present : No
Backend SAS Address 0 : 0x5000C50000...
Backend SAS Address 2 : 0x5000C50000...
BIOS Version : 06.10.00.00
Firmware Build Time :
Product Name : SAS3440X
SubVendor ID : 0x1000
Cache Flush Interval : 0 sec
Device Port Count : 4
Host Interface : PCIX
Alarm Present : No
NVRAM Present : Yes
Backend SAS Address 1 : 0x5000C50000...
Backend SAS Address 3 : 0x0
Firmware Version : 01.16.00.00
Error Level
Date / Time
Description
Information ################ Controller ID: 0 Initialization complete on VD 1
Information ################ Controller ID: 0 Fast initialization started on VD 1
Information ################ Controller ID: 0 Created VD 1
Information 2007-02-08, 13:38:59 Successful log on to the server User: Administrator, Client: 192.168.1.11, Access Mode: Full, Client Time: 2007-02-08,13:38:59
Displaying log from server

▼ To Create a RAID Partition and File System (Windows 2003)

After you create a new RAID configuration, you need to create a partition and format the partition using Windows Disk Manager.

1. On the taskbar, click Start, and then click Run.

2. **In the Open field, enter** `diskmgmt.msc`**, and then click OK.**

   The Windows Disk Manager Initialize Wizard displays.

FIGURE 8-11 Windows Disk Manager Initialize Wizard

Disk Management
File Action View Help
Volume
(C:)
Disk 0
Basic
68.36 GB
Online
Disk 1
Unknown
135.96 GB
Not Initialized
Unallocated
Initialize and Convert Disk Wizard
Welcome to the Initialize and Convert Disk Wizard
This wizard helps you to initialize new disks and to convert empty basic disks to dynamic disks.
You can use dynamic disks to create software-based volumes that can be mirrored, or they can be striped or spanned across multiple disks. You can also expand single-disk and spanned volumes without having to restart the computer.
After you convert a disk to dynamic, you can only use Windows 2000 and later versions of Windows on any volume of that disk.
To continue, click Next.
< Back
Next >
Cancel

3. **Click Cancel to exit the Windows Disk Manager Initialize wizard.**

4. **Right-click the Disk containing the new RAID created by MSM, and then click Initialize Disk.**

FIGURE 8-12 Windows Disk Manager Initialize Disk

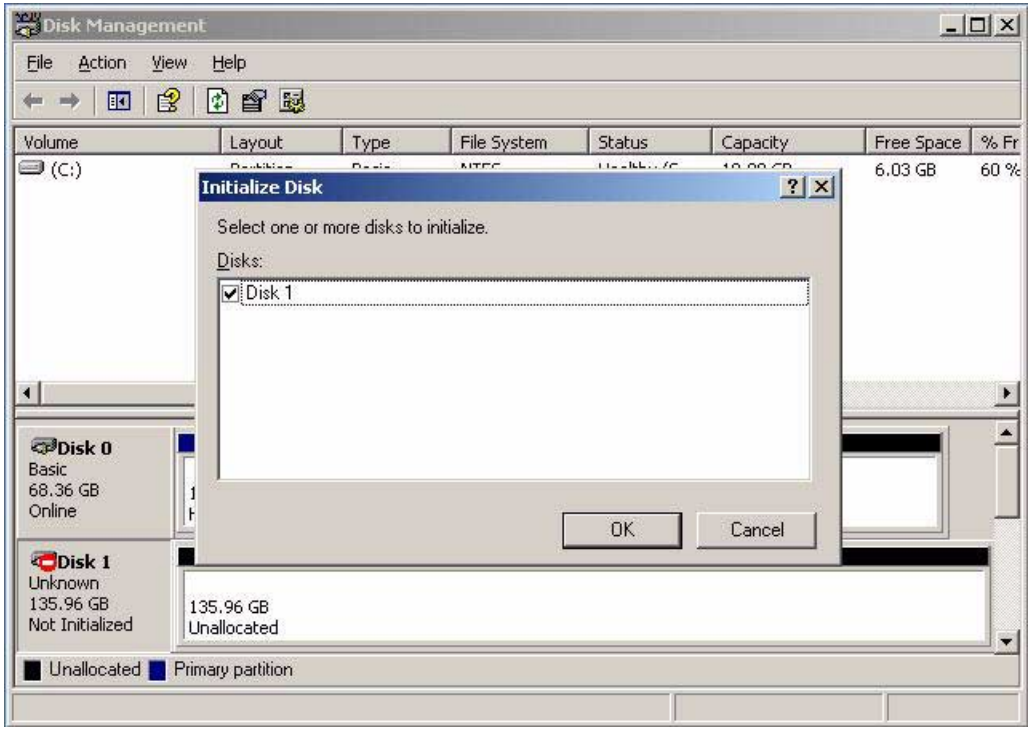

5. **Right-click on the Un-allocated Drive section of the new initialized partition, and then select New Partition.**

FIGURE 8-13 Windows Disk Manager New Partition Wizard

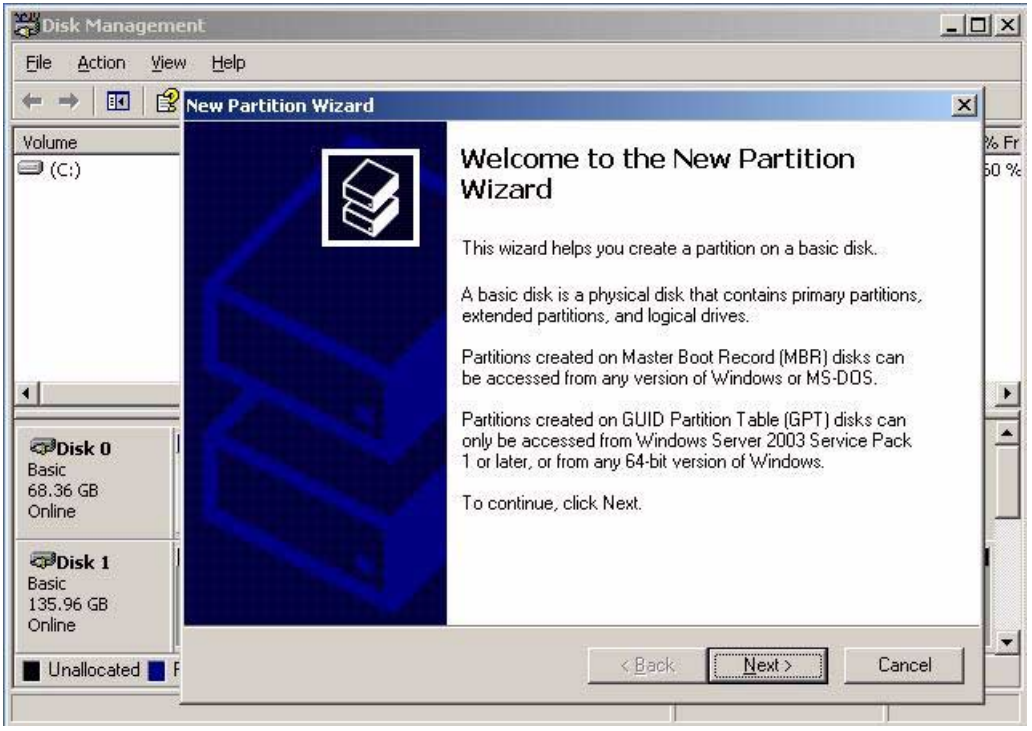
Disk Management
File Action View Help
Volume
(C:)
Disk 0
Basic
68.36 GB
Online
Disk 1
Basic
135.96 GB
Online
Unallocated
New Partition Wizard
Welcome to the New Partition Wizard
This wizard helps you create a partition on a basic disk.
A basic disk is a physical disk that contains primary partitions, extended partitions, and logical drives.
Partitions created on Master Boot Record (MBR) disks can be accessed from any version of Windows or MS-DOS.
Partitions created on GUID Partition Table (GPT) disks can only be accessed from Windows Server 2003 Service Pack 1 or later, or from any 64-bit version of Windows.
To continue, click Next.
< Back
Next >
Cancel

6. **Follow the onscreen prompts to page through the New Partition wizard to create and format a new Windows partition.**

   Once the New Partition wizard completes, the RAID file system is built and available for use. See FIGURE 8-14.

FIGURE 8-14 Windows Display of New RAID File System

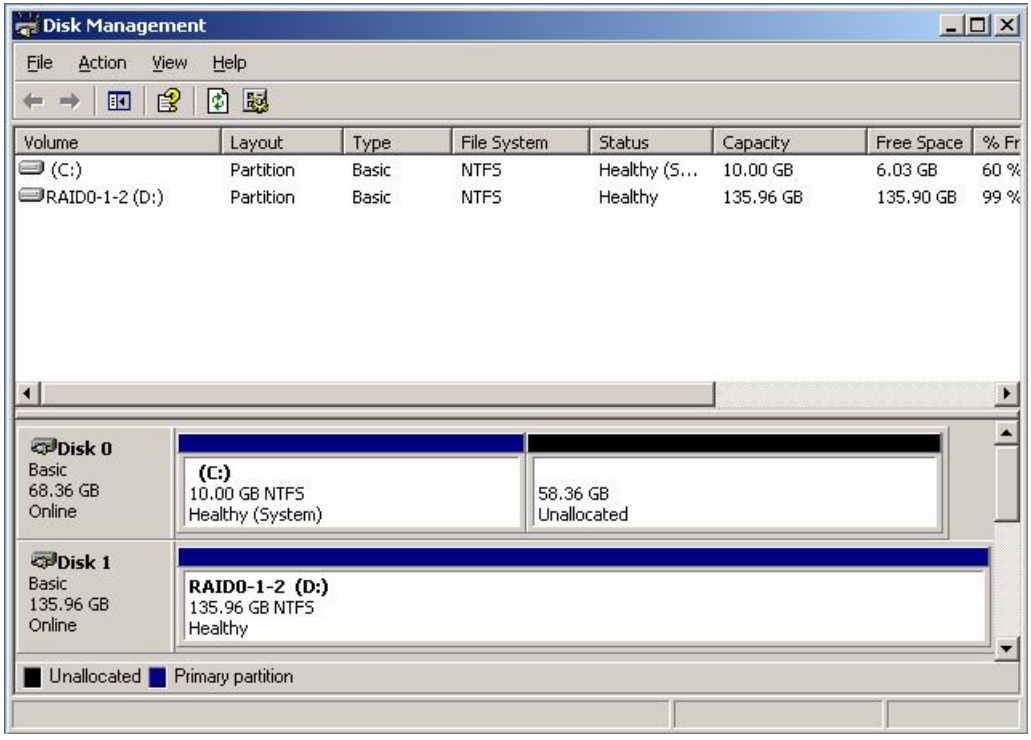

Note – MSM might also be installed from the BIOS. For details see the Service Manual for your product.

Changing Virtual Disk Properties

This section describes the virtual disk parameters that you can set when you use the Guided Configuration or Manual Configuration modes of the Configuration Wizard. You do not necessarily need to change these parameters when you create a storage configuration; you can leave the values at the default settings.

- **Disk Cache Policy:** Select a cache setting for this disk: *Unchanged*, *Enabled*, or *Disabled*.

- **Init State:**
  - No Initialization: The new configuration is not initialized and the existing data on the disks is not overwritten.
  - Fast Initialization: MSM quickly writes zeros to the first and last 10 MB regions of the new virtual disk, and then completes the initialization in the background. This enables you to start writing data to the virtual disk immediately.
  - Full Initialization: A complete initialization is done on the new configuration. You cannot write data to the new virtual disk until the initialization is complete. This might take a long time if the disks are large.

Caution – To prevent an inoperable system, do not initialize a drive on which the operating system is installed. Initializing any drive overwrites the drive contents, including the operating system if it is installed.

You can change a virtual disk's Read Policy, Write Policy, and other properties at any time after the virtual disk is created.

Note – Support is provided for enabling/disabling SMART and Write Cache Enable on physical disks that are not part of a virtual disk, but are connected to a SAS IR controller. This is different from the way in which properties are changed for virtual disks.

▼ To Change a Virtual Disk Property

1. **Select a virtual disk icon in the left panel of the MSM window.**
2. **In the right panel, select the Properties tab, and select Set Virtual Disk Properties.**

   A list of virtual disk properties displays in the right panel.
3. **Change the virtual disk properties as needed in the right panel.**
4. **Click Go to accept the changes.**

Deleting a Virtual Disk

You can delete virtual disks for a number of reasons, including rearranging the storage space.

Caution – To prevent any data loss, including the operating system (OS), back up the data on the virtual disk before you delete it. Ensure the operating system is not installed on the virtual disk you intend to delete.

Note – Virtual drives with a bootable partition cannot be deleted. This prevents you from accidentally deleting a drive that contains the operating system. To delete the virtual drive, you must clear the bootable flag in the partition or remove the partition.

▼ To Delete A Virtual Disk

1. **Back up all user data on the virtual disk you intend to delete.**
2. **In the left panel of the MSM window, select the Logical tab and click the icon of the virtual disk to delete.**
3. **In the right panel, select the Operations tab and select Delete Virtual Disk.**
4. **Click Go.**
5. **When the warning message appears, click Yes to confirm the deletion of the virtual disk.**

Saving a Storage Configuration to Disk

There is a `Save Configuration` command on the **Operations > Advanced Operations > Configuration** menu. The purpose of this command is to save an existing controller configuration to a file so that it can be applied to the same or another controller at a later date.

Caution – This `Save Configuration` command is problematical and is not supported.

Clearing a Storage Configuration From a Controller

To create a new configuration on the controller, the existing storage configuration must be cleared.

Caution – To prevent any data loss, back up the data on the virtual disk connected to the controller before you clear a configuration. Clearing a configuration deletes all data from the disks of the existing configuration.

▼ To Clear a Configuration from a Controller

1. **Select a controller icon in the left panel of the MSM window.**
2. **On the menu bar, select Operations > Advanced Operations > Configuration > Clear Configuration.**

   A warning message displays.
3. **To clear the configuration, click Yes.**

Monitoring System Events and Storage Devices

MSM enables you to monitor the status of disk drives, virtual disks, and other storage devices. This section has the following subsections:

- "Monitoring System Events" on page 90
- "Monitoring Controllers" on page 91
- "Monitoring Disk Drives" on page 92
- "Monitoring Virtual Disks" on page 93
- "Monitoring Rebuilds and Other Processes" on page 94

Monitoring System Events

MSM monitors the activity and performance of all controllers in the system and the storage devices connected to the controllers. When an event occurs (such as the creation of a new virtual disk or the removal of a physical disk drive), an event message appears in the log displayed at the bottom of the MSM window.

Each event in the log includes an error level (Info, Warning, Caution, Fatal, or Dead) a date and timestamp, and a brief description.

Note – Some events in MSM_IR do not display an accurate date and timestamp value in the MSM log. When this is the case, the date and timestamp line will display "#####" as the value.

For example, creating a new RAID 0 or RAID 1 will generate "#####" in the MSM date & timestamp log. Also, swapping hard disks will only display "###################" in the MSM log section.

The Log menu has three options:

- **Save Log**: Saves the current log to a file.
- **Clear Log**: Clears the current log information.
- **Load Log**: Enables you to load a different log file.

Monitoring Controllers

MSM enables you to see the status of all

controllers in the left panel of the MSM window. The controller's status is indicated by the controller icon(s).

- The controller icon by itself indicates that it is operating normally.

- 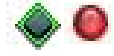

  A red LED next to the icon indicates that the controller has failed. (See "Physical/Logical View Panel" on page 74 for a complete list of device icons.)

▼ To Display All Controller Information

1. **Click a controller icon in the left panel of the MSM window.**
2. **Click the Properties tab in the right panel.**

   The Properties tab lists information such as Product Name, Serial number, Vendor ID, and Host Port Count.

All controller properties are defined in "Glossary" on page 127.

Monitoring Disk Drives

MSM enables you to see the status of all physical disk drives in the left panel of the MSM window. The physical disk drive status is indicated by the physical disk drive icon(s):

-

  The physical disk drive icon by itself indicates that it is operating normally.

- 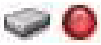

  A red LED next to the icon indicates that the physical disk drive has failed.

▼ To Display Complete Disk Drive Information

1. **Click a disk drive icon in the left panel of the MSM window.**
2. **Click the Properties tab in the right panel.**

   The Properties tab lists information such as Vendor Name, Device ID, and Physical Drive State. All disk drive properties are defined in "Glossary" on page 127.

▼ To Display a Graphical View of a Disk Drive

1. **Click a drive icon in the left panel of the MSM window.**
2. **Click the Graphical View tab.**

In Graphical View, the drive's storage capacity is color-coded according to the legend shown on the screen:

- Blue = Configured space
- White = Available space
- Red = Reserved space

- Green = The space used by the virtual disk when you select a virtual disk from the drop-down menu.

Monitoring Virtual Disks

MSM enables you to see the status of all virtual disks. The virtual disk status is indicated by the virtual disk icon.

-

  The virtual disk icon by itself indicates that it is operating normally.

-

  A yellow LED next to the icon indicates that the virtual disk is in degraded mode.

  For example, if a physical disk has failed, the virtual disk icon reflects this degraded condition

- 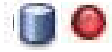

  A red LED next to the icon indicates that the virtual disk has failed. (See "Physical/Logical View Panel" on page 74 for a complete list of device icons.)

▼ To Display a Graphical View of a Virtual Disk

1. **Click a virtual disk icon in the left panel of the MSM window.**
2. **Click the Graphical View tab.**

In Graphical View, the icon for this disk group (array) indicates the virtual disk usage.

- Blue = Indicates how much of the disk group capacity is used by this virtual disk.
- White = Indicates that some of the virtual disk capacity is used by another virtual disk.

Monitoring Rebuilds and Other Processes

MSM enables you to monitor the progress of rebuilds and other lengthy processes in the Group Show Progress window. Open this window by selecting **Group Operations > Show Progress** on the menu bar.

Operations on virtual disks appear in the left panel of the Group Show Progress window, and operations on physical disk drives appear in the right panel. The following operations appear in this window.

- Background initialization of a virtual disk
- Rebuild (see "To Rebuild a Drive on a SAS IR System" on page 95)
- Consistency check (SAS IR controllers only)

To abort any ongoing process:

- Click the **Abort** button next to the status indicator.
- Click **Abort All** to abort all ongoing processes.
- Click **Close** to close the window.

Maintaining and Managing Storage Configurations

This section explains how to use MSM to maintain and manage storage configurations. This section has the following subsections:

- "To Scan for New Drives" on page 94
- "Rebuilding a Drive" on page 95
- "Putting a Drive Offline or Missing" on page 96

▼ To Scan for New Drives

MSM normally detects newly installed disk drives and displays icons for them in the MSM window. If MSM fails to detect a new drive (or drives), you can use the Scan for Foreign Config command to find it.

1. **Select a controller icon in the left panel of the MSM window.**
2. **Select Operations > Scan for Foreign Config.**

   If MSM detects any new disk drives, it displays a list of them on the screen.

3. **Follow the instructions on the screen to complete the disk detection.**

Rebuilding a Drive

Due to the definition of a RAID 1 virtual drive, the system is protected from data loss if a single drive fails. However, the failed drive must be replaced, and the failed drive's data must be rebuilt on a new drive to restore the system to a fault-tolerant condition. You can also rebuild a failed drive, if the failed drive is still operational.

-

  A yellow LED next to the virtual disk indicates that it is in a degraded state; the data is still safe, but data could be lost if another drive fails.

  If you see that the virtual disk is in a degraded state, then view the physical disk in the virtual disk configuration for drive indications.

- 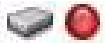

  A red LED next to the drive icon indicates that the drive has failed.

▼ To Rebuild a Drive on a SAS IR System

1. **Record the number of the failed drive(s) (0, 1, 2, 3) in the MSM window.**
2. **Shut down the system, disconnect the power cord, and open the server chassis.**
3. **Find the failed disk drive and remove it from the server chassis.**

   You can identify the disk drive by reading the number (0, 1, 2, 3) on the drive cable. This corresponds to the drive number displayed in the MSM window. Also, the drive 0 cable is color-coded. For an Integrated RAID controller, the hard drive number is on the motherboard next to the cable connector.
4. **Replace the failed disk drive with a new drive of equal or greater capacity.**
5. **Close the computer case, reconnect the power cord, and restart the server.**
6. **Restart MSM.**

   When the new drive spins up, the drive icon changes back to normal status, and the rebuild process begins automatically.

Putting a Drive Offline or Missing

Note – The option to mark a physical disk (drive) as missing does not appear in some versions of MSM. In those versions where "**Mark physical disk as missing**" is disabled or not available, you will see "Mark drive online" and "Rebuild" options.

If a disk drive is currently part of a redundant configuration and you want to use it in another configuration, you can use MSM commands to remove the disk drive from the first configuration for this purpose. When you do this, *all data on that drive is lost*. You can remove the disk drive from the configuration without harming the data on the virtual disk.

Note – If a disk drive in a virtual disk has failed, the drive goes offline. If this happens, you must remove the disk drive and replace it. You cannot make the drive usable for another configuration by using the **Mark physical disk as missing** and **Rescan** commands.

▼ To Put a Drive Offline or Missing

1. **In the left panel of the MSM window, right-click the icon of a disk drive in a redundant virtual disk.**
2. **Select Make drive offline from the pop-up menu.**

   The disk drive status changes to Offline.
3. **Right-click the disk drive icon again and select Mark physical disk as missing.**
4. **Select File > Rescan.**

   The disk drive status changes to Unconfigured Good. At this point, the data on this disk drive is no longer valid.

Known Issues

The following section lists known issues by product.

Sun Fire X4100 M2 /X4200 M2 Server Issues

(Windows 2003 Server) MSM-IR 1.19 Does Not Reflect Correct Disk Count Information (CR 6514389)

On Windows Server 2003 (32- and 64-bit), MSM-IR 1.19 does not show the correct disk count information at the physical and logical level before the RAID set is created. There is no workaround, but this was fixed in software release 1.1 (February, 2007).

Sun Fire X4600 M2 Server Issues

MSM-IR 1.19 Does Not Show Disk Removal Status Correctly in a Non-RAID Configuration (CR 6525255)

In a non-RAID configuration only, the status log does not show an entry when a disk is removed. There is no workaround.

MSM Server and Client Must Be in Same Subnet (6533271)

For firmware prior to Phase 11, MSM 2.29, the MSM server and client must be in the same subnet; otherwise, the applications reports “No server found.”

Sun Blade X6240, X6440 Server Issues

Locate Virtual Disk Function Does Not Light LEDs on Disks Controlled by Server Blade (CR 6732326)

When you have a virtual disk highlighted in the logical view and you choose the Locate Virtual Disk command in the Operations tab, LEDs should blink on all the physical disks that make up the virtual disk. These LEDs are not blinking on disks controlled by the X6240, X6440 server blades.

CHAPTER 9

LSI SNMP Utility

The LSI (SAS IR) SNMP utility is used over SAS connections to monitor MSM activity from a remote station for Windows Server 2003 systems and Linux systems. The LSI SNMP agent requires installation of SNMP service on the server side followed by installation of the LSI SNMP agent on the remote station.

This chapter includes the following sections:

- "Installing and Configuring SNMP Service" on page 99
- "Installing LSI SNMP on a Remote Station" on page 101

Installing and Configuring SNMP Service

This section includes the following topics:

- "To Install the SNMP Service on the Server Side on Windows Server 2003" on page 100
- "To Configure the SNMP Service on the Server Side on Windows Server 2003" on page 100
- "Installing and Configuring the SNMP Server on the Server Side on Linux" on page 101

Note – You must install and configure SNMP service before installing the LSI (SAS IR) SNMP Agent.

▼ To Install the SNMP Service on the Server Side on Windows Server 2003

Note – To complete this procedure, you will need the Windows Server 2003 CD that came with your system.

If the SNMP service has already been installed and configured, skip this procedure.

1. **Open the Control Panel and click Add/Remove Programs.**

   The Add/Remove Programs window displays.

2. **Click Add/Remove Windows Components in the left side of the Add/Remove Programs window.**
3. **Select Management and Monitoring Tools.**
4. **Click Next.**
5. **Insert the Windows Server 2003 CD when prompted.**

The Windows Server 2003 OS extracts the necessary SNMP files and installs the SNMP service on your server.

▼ To Configure the SNMP Service on the Server Side on Windows Server 2003

1. **Open the Control Panel and click Administrative Tools.**
2. **Click Services in Administrative Tools window.**
3. **Select SNMP Service in the Services window.**
4. **Open SNMP Service.**
5. **Click the Security tab and select the** *Accept SNMP Packets From Any Host***.**
6. **If you want to send traps to a host IP, click the Traps tab and select from the list of host IPs.**

Installing and Configuring the SNMP Server on the Server Side on Linux

Linux comes with a SNMP server. However, the RPM included in the LSI package needs to be installed on the system that is to be monitored, that is, the system with LSI Raid.

▼ To Install the SNMP Agent

1. **Unzip** `SAS_IR_SNMP_Linux_Installer-3.xx.xxx.zip`.
2. **Uncompress** `sas_ir_snmp.tar.gz`.

Note – On SLES10, before running the following command, copy `/etc/snmp/snmpd.conf` to `/etc/snmpd.conf`. After running the command, copy `/etc/snmpd.conf` back to `/etc/snmp/snmpd.conf`.

---

3. **Run** `rpm -ivh sas_ir_snmp-3.09.0000.i386.rpm`.
4. **To allow a client machine to run SNMP queries on the server, modify the** `snmpd.conf` **file by adding this line:**

   ```
   com2sec   snmpclient  1.1.1.1   public
   ```

   where 1.1.1.1 is the IP address of the client from which SNMP queries are sent.
5. **To configure the server to send traps to remote clients automatically, add the following line to** `/etc/lsi_mrdsnmp/sas-ir/sas_ir_TrapDestination.conf:`

   1.1.1.1 public

   where 1.1.1.1 is the IP address of the machine you wish to send traps to.

Installing LSI SNMP on a Remote Station

The LSI (SAS IR) SNMP Agent is installed on the remote station to access the server with MSM and Windows Server 2003 installed.

This section has the following topics:

- "To Download the LSI SNMP Agent Files" on page 102

■ "To Install SNMP Files on a Remote Station" on page 102

▼ To Download the LSI SNMP Agent Files

This procedure describes how to download the necessary files for the LSI (SAS IR) SNMP Agent for Windows Server 2003 systems.

Note – Before you install the SAS IR Agent, verify that the SNMP service is already installed and configured on the system. If the SNMP Service is not installed on your system, refer to "Installing and Configuring SNMP Service" on page 99.

To install the SNMP files on the remote station:

1. **Find the following files:**
   - (Windows) `SAS-IR_SNMP_Win_Installer-3.xx-xxxx.zip`

Note – This file is not on the CD or in the online downloads. To obtain it contact your Sun service representative. (Refer to CR 6578969.)

   - (Linux) `SAS-IR_SNMP_Linux_Installer-3.xx-xxxx.zip` is in `/linux/tools/raid`
   - `SUN-PLATFORM-MIB.mib` is in `/common/snmp/`
2. **Place the `SUN-PLATFORM-MIB.mib` file in a folder on the remote station, where you can compile the file.**
3. **Place the `SAS-IR_SNMP_Win_Installer-3.11-0000.zip` or the `SAS-IR_SNMP_Linux_Installer-3.xx-xxxx.zip` file on the remote station in a separate folder where you can extract the files.**

▼ To Install SNMP Files on a Remote Station

1. **Compile the `SUN-PLATFORM-MIB.mib` file in the directory you selected.**
2. **Extract the files from the `SAS-IR_SNMP_Win_Installer-3.11-0000.zip` or the `SAS-IR_SNMP_Linux_Installer-3.xx-xxxx.zip`.**

   The zipped file contains the following files and folders:
   - `DISK1` folder
   - `documents` folder
   - `readme.txt`

3. **Open the `DISK1` folder and run `setup.exe` to install the LSI (SAS IR) SNMP Agent on the remote system.**

4. **Use the SNMP Manager to retrieve the SAS IR data and monitor the MSM activity on the server from the remote station.**

Note – The trap function of SNMP is described in the Integrated Lights Out Management (ILOM) documentation available on your product documentation web site. You will need the MIB file, `LSI-megaRAID_Sas_IR.mib`. This MIB describes the LSI SNMP traps and this MIB must be compiled into the trap-catching utility.

PART IV P A R T IV - Using `raidctl` With Solaris OS

This part describes how to use the Solaris command raidctl to create hardware RAID for any server running the Solaris OS and has the following chapters:

- "Introduction to `raidctl`" on page 107
- "The raidctl Man Page" on page 113

CHAPTER 10

Introduction to `raidctl`

This chapter provides an overview of the LSI Integrated RAID solution for LSI SAS controllers. The chapter includes these sections:

- "What is `raidctl`?" on page 107
- "When to Use `raidctl`" on page 108
- "Using `raidctl` to Create RAID Volumes" on page 108
- "Other Uses For `raidctl`" on page 112

What is `raidctl`?

`raidctl` is a Solaris command that can be used to set up hardware RAID volumes on LSI host bus adapters (HBAs). It is entirely analogous to the LSI BIOS utility described in Part I of this document.

SPARC systems, which run the Solaris OS, do not have a BIOS, so that the LSI BIOS utility is not available to set up RAID on the LSI HBA. `raidctl` takes its place.

`raidctl` can be used to set up RAID volumes for any LSI HBA that can be set up with the LSI BIOS utility. This includes on-board LSI 106x chips and PCI cards based on LSI 106x chips.

Note – Since `raidctl` is a Solaris command, it can be used with *any* server running the Solaris OS, no matter whether it has a SPARC, AMD, or Intel processor. The behavior of `raidctl` is not dependent on which server is running the Solaris OS.

When to Use `raidctl`

Hardware RAID can be set up with `raidctl` before or after the your server's OS is installed. However, if you want to mirror your boot disk, the RAID mirror must be set up *before* OS installation. To do this:

1. **Boot your new server using a remote Solaris OS.**
2. **Use `raidctl` to create your RAID mirror.**
3. **Reboot and install the OS on the mirror.**

Using `raidctl` to Create RAID Volumes

You use the `-c` or `-C` option to create a RAID volume. The `-c` option can only create RAID 0, 1, and 1E volumes. The `-C` option is more general.

The `raidctl -C` and `raidctl -c` commands are described in detail in the raidctl man page, which is reproduced in the next chapter ("The raidctl Man Page" on page 113). Numerous examples are given in the man page.

Disk Names

HBA's might have different connectors to different SCSI buses; these are called channels. In Solaris device file convention, they are represented by the letter `c`, or controller number.

SCSI disks are addressed by target number and logical unit numbers. There could be multiple logical unit numbers up to a maximum of 8 under each target number.

`raidctl` uses two slightly different formats for naming disks:

1. The Solaris canonical format, `c?t?d?`, where `c` is the controller number, `t` is the target number, and `d` is the logical unit number. For example, three disks connected to controller number 2 could be c2t0d0, c2t1d0, and c2t2d0.
2. The `C.ID.L` format, where `C` is the channel number (not the same as the controller number) and `ID` and `L` are once again the target ID and logical unit number.

Obtaining Disk Names in Canonical Format

You can run the format command at the CLI without any parameters to get the names of the available disks in the canonical format. For example,

```
# format
Searching for disks...done
c2t3d0: configured with capacity of 136.71GB
AVAILABLE DISK SELECTIONS:
0. c2t0d0 .........
1. c2t1d0 ..........
2. c2t2d0.........
3. c2t3d0........
#
```

Obtaining Disk Names in `C.ID.L` Format

You can run the `raidctl` command with the `-l` option to obtain the names of the available disks in `C.ID.L` format. For example,

```
# raidctl -l
Controller: 2
Disk: 0.0.0
Disk: 0.1.0
Disk: 0.2.0
Disk: 0.3.0
#
```

Note – Running the `raidctl -l` command also provides the number of the controller, which is 2. This means that the name of these disk in canonical form would be `c2t0d0`, `c2t1d0`, `c2t2d0`, and `c2t3d0` (as was found above by running the `format` command).

The `raictl -c` Command

```
raidctl -c [-f] [-r raid_level] disk1 disk2 [disk3...].
```

Parameter	Description
`-f`	When present, this parameter suppresses warning prompts, such as "Do you really want to..."
`-r raid_level`	`raid_level` can be 0, 1, or 1E. For `raid_level` = 0, there must be two or more disks listed. For `raid_level` = 1, there must be two and only two disks listed. If there are more than two disks listed, `raid_level` must be 0 or 1E. If the `-r` parameter is omitted, `raidctl` will create a RAID 1 volume if there are two disks listed and will fail otherwise.
disk1, disk2, ...	Disk names in Solaris canonical format, `c?t?d?`.

Note – If you run the `raidctl -c` command without the `[-r raid_level]` option, you can only list two disks and you will get a RAID 1 volume. To create a RAID 1E volume, you must list more than two disks *and* you must use the `-r` option.

Here is what happens if you list three disks and do not specify the `raid_level` option:

```
# raidctl -c c2t1d0 c2t2d0 c2t3d0
Creating RAID volume will destroy all data on spare space of member
disks, proceed (yes/no)? yes
Illegal array layout.
#
```

Here is what happens when you do not specify the `raid_level` option, but only list two disks:

```
# raidctl -c c2t1d0 c2t2d0
Creating RAID volume will destroy all data on spare space of member
disks, proceed (yes/no)? y
Volume c2t1d0 is created successfully!
#
```

Although the output did not say so, `c2t1d0 is a RAID 1 volume.`

The `raidctl -C` Command

This command is more general than `raidctl -c` and uses a different format for naming disks (C.ID.L)

```
raidctl -C "disks" [-r raid_level] [-z capacity]
[-s stripe_size] [-f] controller
```

Parameter	Description
`"disks"`	A list of disks in `C.ID.L` format. The list can include disks and sub-volumes, separated by spaces. Sub-volumes are groups of disks separated by spaces but enclosed by parenthesis—for example, (0.0.0 0.1.0).
`-r raid_level`	`raid_level` can be 0, 1, 1E, 5, 10, or 50. See the man page for descriptions of the disks combinations that can be used. If this parameter is omitted, `raidctl` will create a RAID 1 volume if two disks are listed and will fail otherwise. **Note -** The LSI 106x HBA can only form RAID levels 0, 1, and 1E
`-z capacity`	The capacity of the volume that will be created. Can be terabytes, gigabytes, megabytes, etc., entered as 2t, 24g, 256m and so forth. If this parameter is omitted, `raidctl` calculates the maximum volume that can be created from the disks listed.
`-s stripe_size`	Stripe size of the volume that will be created. See the man page for the possible values. If this parameter is omitted, `raidctl` chooses an appropriate value, often 64k.
`-f`	When present, this parameter suppresses warning prompts, such as "Do you really want to..."
controller	Specifies to which HBA (RAID controller) the disks belong. `raidctl -l` will return the controller's ID.

Note – As with `raidctl -c`, you must use the `[-r raid_level]` option unless you are forming a RAID 1 volume with just two disks.

Other Uses For `raidctl`

Depending on the option selected, `raidctl` can also be used to:

- Delete RAID volumes
- Create or delete hot-spares
- Update controller firmware
- Display information about a specified disk
- Display information about a specified volume
- Display information about a specified controller
- Take a snapshot of the RAID configuration

These options are all described in the next chapter, which lists the `raidctl` man page.

CHAPTER 11

The `raidctl` Man Page

This chapter provides a listing of the **raidctl** man page.

The Man Page for `raidctl`

man raidctl

System Administration Commands raidctl(1M)

NAME

raidctl - RAID hardware utility

SYNOPSIS

raidctl -C “disks” [-r raid_level] [-z capacity] [-s stripe_size] [-f] controller

raidctl -d [-f] volume

raidctl -F filename [-f] controller...

raidctl -a {set | unset} -g disk {volume | controller}

raidctl -p “param=value” [-f] volume

raidctl -c [-f] [-r raid_level] disk1 disk2 [disk3...]

raidctl -l -g disk controller

raidctl -l volume

raidctl -l controller...

raidctl [-l]

raidctl -S [volume | controller]

raidctl -S -g disk controller

raidctl -h

DESCRIPTION

The raidctl utility is a hardware RAID configuration tool that supports different RAID controllers by providing a CLI (command-line interface) to end-users to create, delete or display RAID volume(s). The utility can also used to set properties of a volume, assign hot-spare (HSP) disks to volumes or controllers, and to update firmware/code/BIOS for RAID controllers.

The raidctl utility requires privileges that are controlled by the underlying file-system permissions. Only privileged users can manipulate the RAID system configuration. If a non-privileged user attempts to run `raidctl`, the command fails with an exit status of 1.

The raidctl utility defines a set of command line options to provide management for full feature RAID controllers. Since the supported features may vary with different RAID controllers, not all options are supported for a given RAID controller. The user can use raidctl to list the type of a given controller and the firmware version to determine the supported features.

Currently, raidctl supports the following RAID controllers:

- LSI1020, LSI1030, LSI1064, and LSI1068 SCSI HBAs, that are maintained by the mpt driver, on X86-32/64 and SPARC platforms.

OPTIONS

The following options are supported:

-C “disks” [-r raid_level] [-z capacity] [-s stripe_size] [-f] controller

Create a RAID volume using specified disks.

When creating a RAID volume using this option, the identity of the newly created volume is automatically generated and raidctl reports it to the user.

The argument specified by this option contains the elements used to form the volume that will be created. Elements can be either disks or sub-volumes, where disks are separated by space(s) and a sub-volume is a set of disks grouped by parenthesis. All disks should be in C.ID.L expression (for example, 0.1.2 represents a physical disk of channel 0, target id 1, and logical unit number 2). The argument must match the RAID level specified by the -r option, even if it’s omitted. This means the argument can only be:

for RAID 0	At least 2 disks
for RAID 1	Only 2 disks
for RAID 1E	At least 3 disks
for RAID 5	At least 3 disks
for RAID 10	At least 2 sub-volumes, each sub-volume must be formed by 2 disks
for RAID 50	At least 2 sub-volumes, each sub-volume must be formed by at least 3 disks, and the disk amount in each sub-volume should be the same

For example, the expression “0.0.0 0.1.0” means that the 2 specified disks form a RAID volume, which can either be a RAID 0 or a RAID 1 volume. “(0.0.0 0.1.0)(0.2.0 0.3.0)” means that the first 2 disks and the last 2 disks form 2 sub-volumes, and that these 2 sub-volumes form a RAID 10 volume. See the EXAMPLES section for more samples.

The -r option specifies the RAID level of the volume that will be created. Possible levels are 0, 1, 1E, 5, 10, 50. If this option is omitted and only two disks are listed, `raidctl` creates a RAID 1 volume by default—otherwise, it fails.

The -z option specifies the capacity of the volume that will be created. The unit can be tera-bytes, giga-bytes, or mega-bytes (for example, 2t, 10g, 20m, and so on). If this option is omitted, raidctl calculates the maximum capacity of the volume that can be created by the specified disks and uses this value to create the volume.

The -s option specifies the stripe size of the volume that will be created. The possible values are 512, 1k, 2k, 4k, 8k, 16k, 32k, 64k, or 128k. If this option is omitted, raidctl chooses an appropriate value for the volume (for example, 64k).

In some cases, the creation of a RAID volume may cause data on specified disks to be lost (for instance, on LSI1020, LSI1030, LSI1064, or LSI1068 HBAs), and raidctl prompts the user for confirmation about the creation. Use the -f option to force the volume creation without prompting the user for confirmation.

The controller argument is used to identify which RAID controller the specified disks belongs. The -l option can be used to list the controller’s ID number.

-d [-f] volume

Delete the RAID volume specified as volume. The volume is specified in canonical form (for example, c0t0d0).

When a volume is deleted, all data is lost. Therefore, unless the -f option is specified, raidctl prompts the user for confirmation before deleting the volume.

When a RAID 1 volume is deleted from a LSI1020, LSI1030, LSI1064, or LSI1068 HBA, the primary and secondary disks are “split”. If the volume was in SYNCING state, the primary will contain the data, and the secondary will not. If the volume state was OPTIMAL, both disks will contain a complete image of the data.

-F filename [-f] controller...

Update the firmware running on the specified controller(s). The raidctl utility prompts the user for confirmation of this action, unless the -f option is provided.

-a {set | unset} -g disk {volume | controller}

If the volume is specified, raidctl sets or unseats the disk as a local hot-spare disk dedicated to the volume, depending on the value specified by the -a option. If the controller is specified, raidctl sets or unseats the disk as a global hot-spare disk.

-p “param=value” [-f] volume

Change the property value for a given RAID volume. Currently, only the cache write policy is changeable (“on” or “off”). So Purim can only be the string “up” (SET_WR_POLICY), and value can be either “on” or “off”.

Changing a RAID volume’s property may affect the internal behavior of the RAID controller, so raidctl prompts the user for a confirmation before applying the change, unless the -f option is specified.

-c [-f] [-r raid_level] disk1 disk2 [disk3...]

Create a volume using the specified disks. This is an alternative to the -C option with similar functionality. This option is preserved for compatibility reasons, but only works with LSI1020, LSI1030, LSI1064, and LSI1068 HBAs to create RAID 0, RAID 1, or RAID 1E volumes. For other HBAs, the user can only use the -C option.

The -r option can be used to specify the RAID level of the target volume. If this option is omitted and only two disks are listed, `raidctl` creates a RAID 1 volume by default—otherwise, it fails.

Disks must be specified in Solaris canonical format (for example, c0t0d0).

Creating a RAID 1 volume with this option replaces the contents of disk2 with the contents of disk1.

When the user creates a RAID volume with this option, the RAID volume assumes the identity of disk1. Other disks become invisible and the RAID volume appears as one disk.

Creating a volume with this option is by default interactive. The user must answer a prompt affirmatively to create the volume. Use the -f option to force the volume creation without prompting the user for confirmation.

-l -g disk controller

Display information about the specified disk of the given controller. The output includes the following information:

Disk	Displays the disk in C.ID.L expression disk.
Vendor	Displays the vendor ID string.
Product	Displays the product ID string.
Capacity	Displays the total capacity of the disk.
Status	Displays the current status of disk. The status can be either "GOOD" (operating normally) or "FAILED" (non-functional).
HSP	Indicates if the disk has been set as a global hot-spare disk, local hot-spare disk, or a normal one. If it is a local hot-spare disk, all volumes which this disk is assigned to are displayed.

-l volume

Display information about the specified volume. The output includes the following information:

Volume	Displays volume in canonical format.
Sub	Displays sub-volumes, if the specified volume is of RAID 10 or RAID 50 volume.
Disk	Displays all disks that form the specified volume.
Stripe Size	Displays the stripe size of the volume.

Status	Displays the status of the specified volume, or the sub-volumes/disks than form the specified volume. For a volume, the status can be “OPTIMAL” (operating optimally), “DEGRADED” (operating with reduced functionality), “FAILED” (non-functional), or “SYNC” (disks are syncing). For a disk, the status can be “GOOD” or FAILED”.
Cache	Indicates whether the cache is applied to I/O write activities. The cache can be either “ON” or “OFF”.
RAID level	Displays the RAID level. The RAID level can be either 0, 1, 1E, 5, 10, or 50.

-l controller...

Display information about the specified controller(s). The output includes the following information:

Controller	Displays the RAID controller’s ID number.
Type	Displays the RAID controller’s product type.
fw_version	Displays the controller’s firmware version.

[-l]

List all RAID related objects that the raidctl utility can manipulate, including all available RAID controllers, RAID volumes, and physical disks. The -l option can be omitted.

The output includes the following information

:

Controller	Displays the RAID controller’s ID number.
Volume	Displays the logical RAID volume name.
Disk	Displays the RAID disk in C.ID.L expression.

-S [volume | controller]

Takes a snapshot of the RAID configuration information including all available RAID devices, RAID controllers, volumes, and disks.

Each line of the output specifies a RAID device and its related information, separated by space(s). All volumes and disks belong to the last specified controller.

The output lists the following information:

Controller	Displays the controller ID number, and the controller type string in double-quotation marks.
Volume	Displays the RAID volume name, number of component disks, the C.ID.L expression of the component disks, the RAID level, and the status. The status can be either OPTIMAL, DEGRADED, FAILED, or SYNCING.
Disk	Displays the C.ID.L expression of the disk, and the status. The status can be either GOOD, FAILED, or HSP (disk has been set as a stand-by disk).

If a volume or a controller is specified, a snapshot is only taken of the information for the specified volume or controller.

-S -g disk controller

Takes a snapshot of the information for the specified disk.

-h

Print out the usage string.

EXAMPLES

Example 1 Creating the RAID Configuration

The following command creates a RAID 0 volume of 10G on controller 0, and the stripe size will be set to 64k:

```
# raidctl -C "0.0.0 0.2.0" -r 0 -z 10g -s 64k 0
```

The following command creates a RAID 1 volume on controller 2:

```
# raidctl -C "0.0.0 1.1.0" -r 1 2
```

The following command creates a RAID 5 volume on controller 2:

```
# raidctl -C "0.0.0 0.1.0 0.2.0" -r 5 2
```

The following command creates a RAID 10 volume on controller 0:

```
# raidctl -C "(0.0.0 0.1.0)(0.2.0 0.3.0)" -r 10 0
```

The following command creates a RAID 50 volume on controller 0:

```
# raidctl -C "(0.0.0 0.1.0 0.2.0)(0.3.0 0.4.0 0.5.0)" -r 50 0
```

Example 2 Displaying the RAID Configuration

The following command displays all available controllers, volumes, and disks:

```
# raidctl -l
Controller: 0
Controller: 2
    Volume:c2t0d0
    Disk: 0.0.0
    Disk: 0.1.0
    Disk: 0.2.0
    Disk: 0.3.0(HSP)
```

The following command displays information about controller 2:

```
# raidctl -l 2
```

Controller	Type	Fw_version
c2	LSI 1030	1.03.39.00

The following command displays information about the specified volume:

```
# raidctl -l c2t0d0
```

Volume	Sub	Disk	Size	Stripe Size	Status	Cache	RAID Level
c2t0d0			10240M	64K	OPTIMAL	ON	RAID5
		0.0.0	5120M		GOOD		
		01.0	5120M		GOOD		
		0.2.0	5120M		GOOD		

The following command displays information about disk 0.3.0 on controller 2:

```
# raidctl -l -g 0.3.0 2
```

Disk	Vendor	Product	Capacity	Status	HSP
0.3.0	MAXTOR	ATLAS10K4_36SC	34732M	GOOD	c2t0d0

Example 3 Deleting the RAID Configuration

The following command deletes a volume:

```
# raidctl -d c0t0d0
```

Example 4 Updating Flash Images on the Controller

The following command updates flash images on the controller 0:

```
# raidctl -F lsi_image.fw 0
```

Example 5 Setting or unsetting a hot-spare disk

The following command sets disk 0.3.0 on controller 2 as a global hot-spare disk:

```
# raidctl -a set -g 0.3.0 2
```

The following command sets disk 0.3.0 on controller 2 as a local hot-spare disk to volume c2t0d0:

```
# raidctl -a set -g 0.3.0 c2t0d0
```

The following command converts disk 0.3.0 on controller 2 from a global hot-spare disk to a normal one:

```
# raidctl -a unset -g 0.3.0 2
```

The following command removes disk 0.3.0 from being a local hot-spare disk from volume c2t0d0:

```
# raidctl -a unset -g 0.3.0 c2t0d0
```

Example 6 Setting the volume's property

The following command sets the write policy of the volume to "off":

```
# raidctl -a set -p "wp=off" c0t0d0
```

Example 7 Creating volumes with the -c option

The following command creates a RAID 1 volume:

```
# raidctl -c c0t0d0 c0t1d0
```

The following command creates a RAID 0 volume:

```
# raidctl -c -r 0 c0t1d0 c0t2d0 c0t3d0
```

Example 8 Taking a snapshot of the RAID configuration information

The following command takes a snapshot of all RAID devices:

```
# raidctl -S
1 "LSI 1030"
c1t1d0 2 0.2.0 0.3.0 1 DEGRADED
0.2.0 GOOD
0.3.0 FAILED
```

The following command takes a snapshot about volume c1t0d0:

```
# raidctl -S c1t0d0
c1t0d0 2 0.0.0 0.1.0 1 OPTIMAL
```

The following command takes a snapshot about disk 0.1.0 on controller 1:

```
# raidctl -S -g 0.1.0 1
0.1.0 GOOD
```

EXIT STATUS

The following exit values are returned:

0 Successful completion.

1 Invalid command line input or permission denied.

2 Request operation failed.

ATTRIBUTES

See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu
Interface Stability	Evolving

SEE ALSO

attributes(5), mpt(7D)

System Administration Guide: Basic Administration

WARNINGS

Do not create raid volumes on internal SAS disks if you are going to use the Solaris Multipathing I/O feature (also known as MPxIO). Creating a new raid volume under Solaris Multipathing will give your root device a new GUID which does not match the GUID for the existing devices. This will cause a boot failure since your root device entry in /etc/vfstab will not match.

PART V Glossary and Index

- "Glossary" on page 127
- "Index" on page 133

Glossary

RAID Terminology

array See *disk group*.

caching The process of using a high speed memory buffer to speed up a computer system's overall read/write performance. The cache can be accessed at a higher speed than a disk subsystem. To improve read performance, the cache usually contains the most recently accessed data, as well as data from adjacent disk sectors. To improve write performance, the cache may temporarily store data in accordance with its write-back policies.

consistency check An operation that verifies that all stripes in a virtual disk with a redundant RAID level are consistent and that automatically fixes any errors. For RAID 1 disk groups, this operation verifies correct mirrored data for each stripe.

controller A chip that controls the transfer of data between the microprocessor and memory or between the microprocessor and a peripheral device such as a physical disk. RAID controllers perform RAID functions such as striping and mirroring to provide data protection. MSM runs on the SAS Integrated RAID controller.

current write policy A virtual disk property that indicates whether the virtual disk currently supports *write back* or *write through* caching mode.

default write policy A virtual disk property indicating whether the default write policy is *Write through* or *Write back*.

See *current write policy* for a definition of these policies.

device driver Software that allows the operating system to control a device such as a printer. Many devices do not work properly unless the correct driver is installed in the computer.

device ID A controller or physical disk property indicating the manufacturer-assigned device ID.

device port count A controller property indicating the number of ports on the controller.

disk cache policy A virtual disk property indicating whether the virtual disk cache is enabled, disabled, or unchanged from its previous setting.

disk group A logical grouping of disks attached to a RAID controller on which one or more virtual disks can be created, such that all virtual disks in the disk group use all of the physical disks in the disk group.

disk subsystem A collection of disks and the hardware that controls them and connects them to one or more controllers. The hardware can include an intelligent controller, or the disks can attach directly to a system I/O bus controller.

fast initialization A mode of initialization that quickly writes zeros to the first and last sectors of the virtual disk. This enables you to start writing data to the virtual disk immediately while the initialization is running in the background.

fault tolerance The capability of the disk subsystem to undergo a single drive failure per disk group without compromising data integrity and processing capability. The SAS Integrated RAID controller provides fault tolerance through redundant in RAID 1 disk groups.

foreign configuration A RAID configuration that already exists on a replacement set of physical disks that you install in a computer system. MSM enables you to import the existing configuration to the RAID controller, or you can clear the configuration so you can create a new one.

formatting The process of writing a specific value to all data fields on a physical disk, to map out unreadable or bad sectors. Because most physical disks are formatted when manufactured, formatting is usually done only if a physical disk generates many media errors.

global hot spare One or two disk drives per controller can be configured as *global hot-spare* disks, to protect data on the IM/IME volumes configured on the controller. If the firmware fails one of the mirrored disks, it automatically replaces the failed disk with a hot-spare disk and then re-synchronizes the mirrored data. The firmware is automatically notified when the failed disk has been replaced, and it then designates the failed disk as the new hot-spare.

hole In MSM (MegaRAID Storage Manager)-IR, a *hole* is a block of empty space in a disk group that can be used to define a virtual disk.

host interface A controller property indicating the type of interface used by the computer host system: for example, *PCIX*.

host port count A controller property indicating the number of host data ports currently in use.

host system Any computer system on which the controller is installed. Mainframes, workstations, and stand-alone desktop systems can all be considered host systems.

hot spare See *global hot-spare*.

initialization The process of writing zeros to the data fields of a virtual disk and, in fault-tolerant RAID levels, generating the corresponding parity to put the virtual disk in a Ready state. Initialization erases all previous data on the physical disks. Disk groups will work without initializing, but they can fail a consistency check because the parity fields have not been generated.

IO policy A virtual disk property indicating whether Cached IO or Direct IO is being used. In Cached IO mode, all reads are buffered in cache memory. In Direct IO mode, reads are not buffered in cache memory. Data is transferred to cache and the host concurrently. If the same data block is read again, it comes from cache memory. (The IO Policy applies to reads on a specific logical drive. It does not affect the Read-ahead cache.)

migration The process of moving virtual disks from one controller to another by disconnecting the physical disks from one controller and attaching them to another one. The firmware on the new controller will detect and retain the virtual disk information on the physical disks.

mirroring The process of providing complete data redundancy with two physical disks by maintaining an exact copy of one disk's data on the second physical disk. If one physical disk fails, the contents of the other physical disk can be used to maintain the integrity of the system and to rebuild the failed physical disk.

name A virtual disk property indicating the user-assigned name of the virtual disk.

offline A physical disk is offline when it is part of a virtual disk but its data is not accessible to the virtual disk.

physical disk A nonvolatile, randomly addressable device for storing data. Physical disks are re-writable and are commonly referred to as disk drives.

physical drive state A physical disk drive property indicating the status of the drive. A physical disk drive can be in one of the following states:

Unconfigured Good: A disk accessible to the RAID controller but not configured as a part of a virtual disk or as a hotspare.

Online: A physical disk can be accessed by the RAID controller and is part of the virtual disk.

Rebuild: A physical disk to which data is being written to restore full redundancy for a virtual disk.

Failed: A physical disk that was originally configured as Online but on which the firmware detects an unrecoverable error.

Unconfigured Bad: A physical disk on which the firmware detects an unrecoverable error; the physical disk was Unconfigured Good or the physical disk could not be initialized.

Missing: A physical disk that was Online, but which has been removed from its location.

Offline: A physical disk that is part of a virtual disk but which has invalid data.

None: A physical disk with the unsupported flag set. An Unconfigured Good or Offline physical disk that has completed the prepare for removal operation.

physical drive type A physical disk drive property indicating the characteristics of the drive.

product info A physical disk property indicating the vendor-assigned model number of the drive.

product name A controller property indicating the manufacturing name of the controller.

RAID A group of multiple, independent disk drives that provide high performance by increasing the number of disks used for saving and accessing data. A RAID disk group improves I/O performance and data availability. The group of disk drives appears to the host system as a single storage unit or as multiple logical disks. Data throughput improves because several physical disks can be accessed simultaneously. RAID configurations also improve data storage availability and fault tolerance.

RAID 0 RAID 0 uses data striping on two or more disk drives to provide high data throughput, especially for large files in an environment that requires no data redundancy.

RAID 1 RAID 1 uses data mirroring on a pair of disk drives so that data written to one physical disk is simultaneously written to the other physical disk. RAID 1 works well for small databases or other small applications that require complete data redundancy.

RAID 1E RAID 1E Integrated Mirroring Enhanced) uses data mirroring on three to ten disks. (Requires Integrated RAID firmware v1.20.00 or above.)

RAID level A virtual disk property indicating the RAID level of the virtual disk.

read policy A controller attribute indicating the current read policy mode. In *Always read ahead* mode, the controller reads sequentially ahead of requested data and stores the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data. In *No read ahead* mode, read-ahead capability is disabled. In *Adaptive read ahead* mode, the controller begins using read-ahead read policy if the two most recent disk accesses occurred in sequential sectors. If the read requests are random, the controller reverts to *No read ahead* mode.

rebuild The regeneration of all data to a replacement disk in a redundant virtual disk after a physical disk failure. A disk rebuild normally occurs without interrupting normal operations on the affected virtual disk, though some degradation of performance of the disk subsystem can occur.

reclaim virtual disk A method of undoing the configuration of a new virtual disk. If you highlight the virtual disk in the Configuration Wizard and click the **Reclaim** button, the individual disk drives are removed from the virtual disk configuration.

redundancy A property of a storage configuration that prevents data from being lost when one physical disk fails in the configuration.

redundant configuration A virtual disk that has redundant data on physical disks in the disk group that can be used to rebuild a failed physical disk. The redundant data can be parity data striped across multiple physical disks in a disk group, or it can be a complete mirrored copy of the data stored on a second physical disk. A redundant configuration protects the data in case a physical disk fails in the configuration.

revision level A physical disk property that indicates the revision level of the disk's firmware.

SAS Serial Attached SCSI. SAS is a serial, point-to-point, enterprise-level device interface that leverages the SCSI protocol set. The SAS interface provides improved performance, simplified cabling, smaller connectors, lower pin count, and lower power requirements when compared to parallel SCSI.

SCSI device type A physical drive property indicating the type of the device, such as *Disk Drive*.

serial no. A controller property indicating the manufacturer-assigned serial number.

stripe size A virtual disk property indicating the data stripe size used in the virtual disk. See *striping*.

striping A technique used to write data across all physical disks in a virtual disk. Each stripe consists of consecutive virtual disk data addresses that are mapped in fixed-size units to each physical disk in the virtual disk using a sequential pattern. For example, if the virtual disk includes five physical disks, the stripe writes data to physical disks 1 through 5 without repeating any of the physical disks. The amount of space consumed by a stripe is the same on each physical disk. Striping by itself does not provide data redundancy.

subvendor ID A controller property that lists additional vendor ID information about the controller.

vendor ID A controller property indicating the vendor-assigned ID number of the controller.

vendor info A physical disk drive property listing the name of the vendor of the drive.

virtual disk (VD) A storage unit created by a RAID controller from one or more physical disks. Although a virtual disk may be created from several physical disks, it is seen by the operating system as a single disk. Depending on the RAID level used, the virtual disk may retain redundant data in case of a disk failure.

virtual disk state A virtual disk property indicating the condition of the virtual disk. Examples include *Optimal* and *Degraded*.

write back caching The controller sends a data transfer completion signal to the host when the controller cache has received all the data in a disk write transaction. Data is written to the disk subsystem in accordance with policies set up by the controller. These policies include the amount of dirty/clean cache lines, the number of cache lines available, and elapsed time from the last cache flush.

write through caching The controller sends a data transfer completion signal to the host when the disk subsystem has received all the data in a transaction.

write policy See *Default Write Policy*.

Index

A

accessing servers
- different subnet, 69, 70
- enter IP address, 70
- multiple severs, 68

C

creating a RAID0 Array, 53
creating a RAID1 array, 56
creating RAID partition, MSM
- Windows, 83

D

deleting a RAID array, 57
Disk Manager initialization wizard, 84

F

failing RAID0, 55

H

hd utility
- installing on Solaris OS, 37

I

Integrated Mirroring (IM), 38
integrated RAID features, 2

L

log in to server, MSM, 70
Logical view, MSM, 74
LSI (SAS IR) SNMP agent, 99
LSI cfggen utility, 37
LSI SNMP
- Agent
  - file installation, 102
  - remote station installation, 101
- SAS IR agent, 99
- server configuration, 100
- server, Windows Server 2003 installation, 100
- utility (SAS IR), 99

M

Microsoft Developers Network (MSDN) web
- site, 63

mirroring
- activating an array, 19
- creating a second IM or IME volume, 15
- creating global hot-spare disks, 16
- creating IM and IME volumes, 12
- deleting an array, 20
- disk write caching, 8
- hot swapping, 7
- hot-spare disk, 8
- IM and IME features, 4
- IM/IME configuration overview, 11
- IM/IME description, 5
- integrated raid firmware, 7
- locating a disk drive or multiple disk drives in a
  - volume, 20
- managing hot-spares, 16
- media verification, 8
- metadata support, 7
- re-synchronization, 7
- selecting a boot disk, 21

mirroring (continued)
SMART support, 8
synchronizing an array, 19
viewing volume properties, 18
write journaling, 8
MSM
accessing different server subnet, 69
Advanced Operations menu, 78
clearing storage configuration, 90
controller icon, 74
creating RAID partition
Windows, 83
degraded state icon, 74
deleting virtual disks, 88
description, 61
detecting newly installed disk drives, 95
device failure icon, 75
disk drive icon, 74
disk group icon, 75
driver installation, 63
File menu, 77
Full Access log in to network, 72
Graphical tab, 76
Group Operations menu, 78
GUI, 61
Help menu, 78
host server icon status, 69
host server login, 70
installation files
Linux, 64
installing on
Linux, 63
Windows Server 2003, 62
LSI SAS106x IR controllers, 61
MegaRAID Storage Manager defined, 61
monitor
disk drive status, 90
disk drives, 92
rebuilding progress, 94
storage devices, 90
virtual disks, 90, 93
multiple servers on screen, 68
New Partition wizard
Windows, 86
Operations menu, 77
Operations tab, 76
physical drive icon, 75
physical/Logical window description, 74
port icon, 74
Properties tab, 76
RAID
configuration wizard, 78
Size selection, 81
terminology, 127
setting bootable partitions, 78
system event log, 75
using RAID configuration wizard, 78
using Shift key to select two drives, 80
View Only, server login, 72
Virtual Disk Creation dialog, 81
virtual disk icon, 75
Windows Disk Manager, 84
Windows Firewall blocking launch, 68
MSM-IR
Log menu, 78
remote station monitoring, 99

N

network login, MSM, 72
New Partition wizard, MSM
Windows, 86

P

Physical view, MSM, 74

R

raidctl
creating raid volumes, 108
disk names, 108
man page, 113
obtaining disk names in C.ID.L format, 109
obtaining disk names in canonical format, 109
other uses for raidctl, 112
raictl -c command, 110
raidctl -C command, 111
what is it?, 107
when to use, 108
rebuilding a RAID1 array, 56
remote station
downloading the LSI SNMP agent files, 101
installing LSI SNMP agent files, 102

S

server only login, MSM, 72

SNMP service
server side
configuration, 100
installation, 100
SNMP utility
see also LSI SNMP
striping
activating a selected array, 32
activating an array, 32
creating a second IS volume, 30
creating IS volumes, 28
deleting an array, 32
disk write caching, 26
integrated striping firmware, 26
introduction, 23
IS configuration overview, 27
IS description, 24
IS features, 24
locating a disk drive or multiple disk drives in a volume, 33
metadata support, 26
selecting a boot disk, 33
SMART support, 26
viewing is volume properties, 31

W

Windows Preinstallation Environment (WinPE), 38